



PUBLIC

SAP HANA Platform 2.0 SPS 03

Document Version: 1.1 – 2018-10-31

SAP HANA Administration Guide

Content

- 1 **SAP HANA Administration Guide.** **10****
- 2 **Administration Information Map.** **11****
- 3 **Database Administration Tasks at a Glance.** **15****
- 4 **SAP HANA System Architecture Overview.** **17****
- 4.1 Server Architecture of Tenant Databases. 17
 - Tenant Databases. 19
 - The System Database. 19
 - Administration of Tenant Databases. 20
- 4.2 Server Architecture of SAP HANA XS Advanced Runtime Platform. 22
- 4.3 Multiple-Host (Distributed) Systems. 25
 - Scale-Out Architecture of Tenant Databases. 26
 - Scale-Out Architecture of SAP HANA XS Advanced Runtime Platform. 28
- 4.4 SAP HANA Services. 30
- 4.5 System Limitations. 35
- 5 **SAP HANA Administration Tools.** **39****
- 5.1 SAP HANA Cockpit. 42
 - Set up SAP HANA Cockpit for the First Time. 43
 - Open SAP HANA Cockpit. 47
 - Authorizations Needed for Monitoring and Administration. 48
 - Open the SAP HANA Database Explorer (SAP HANA Cockpit). 52
 - Setup and Administration with the Cockpit Manager. 59
 - Security Aspects of SAP HANA Cockpit. 86
 - Managing Resources, Users, and Groups with the Cockpit APIs. 98
 - Using XS CLI Commands to Troubleshoot the Cockpit. 112
- 5.2 SAP HANA Studio. 113
 - Open the SAP HANA Administration Console. 114
 - Execute SQL Statements in SAP HANA Studio. 118
 - Managing SAP HANA Systems in SAP HANA Studio. 120
 - SAP HANA Studio Administration Preferences. 138
- 5.3 SAP HANA Hardware Configuration Check Tool for Tailored Data Center Integration. 142
 - Install the SAP HANA Hardware Configuration Check Tool. 143
- 5.4 SAP Solution Manager for SAP HANA Administration. 144
 - Connecting SAP Solution Manager to SAP HANA. 145
 - Central Monitoring and Administration with SAP Solution Manager. 159

	Analyzing the Root Cause of Problems.	160
	Controlling Change.	160
6	System Administration.	162
6.1	Aggregate Monitoring and Administration.	162
	Managing Multiple Resources in SAP HANA Cockpit.	163
	Monitoring Multiple Systems in SAP HANA Studio.	176
6.2	Starting and Stopping SAP HANA Systems.	179
	Starting and Stopping Systems in SAP HANA Cockpit.	179
	Starting and Stopping Systems in SAP HANA Studio.	181
	Starting and Stopping Systems with SAPControl.	187
	Restart Sequence.	188
6.3	Managing Tenant Databases.	189
	Creating and Configuring Tenant Databases.	189
	Monitoring and Managing Tenant Databases.	243
	Configuring Memory and CPU Usage for Tenant Databases.	266
	Using SAP Web Dispatcher for Load Balancing with Tenant Databases.	280
6.4	Configuring SAP HANA System Properties (INI Files).	291
	Database-Specific Configuration Parameters.	293
	Configuring System Properties in SAP HANA Cockpit.	297
	Configuring System Properties in SAP HANA Studio.	301
	Configure System Usage Type.	304
	Reserve Connections for Administrators.	305
6.5	Managing SAP HANA Licenses.	305
	License Keys for the SAP HANA Database.	306
	Managing Licenses in SAP HANA Cockpit.	308
	Managing Licenses in SAP HANA Studio.	314
6.6	Monitoring the SAP HANA Database.	317
	Monitoring in SAP HANA Cockpit.	317
	Monitoring in SAP HANA Studio.	358
	System and Statistics Views.	388
	The Statistics Service.	389
6.7	Managing and Monitoring the Performance of SAP HANA.	394
	Monitoring, Managing, and Analyzing Performance in SAP HANA Cockpit	396
	Monitoring and Analyzing Performance in SAP HANA Studio.	448
	Performance: Using Hints to Query Data Snapshots.	457
	Persistent Data Storage in the SAP HANA Database.	462
	Memory Usage in the SAP HANA Database.	470
6.8	Managing Tables.	479
	Columnar and Row-Based Data Storage.	479
	The JSON Document Store.	484
	Basic Table Management in SAP HANA Cockpit.	486

	Basic Table Management in SAP HANA Studio.	495
	Table and Catalog Consistency Checks.	508
	Memory Management in the Column Store.	518
	The Delta Merge Operation.	526
	Data Compression in the Column Store.	538
	Table Partitioning.	542
	Table Placement.	574
	Table Replication.	582
	Redistributing Tables in a Scaleout SAP HANA System.	600
6.9	Workload Management.	621
	Workload in the Context of SAP HANA.	621
	Controlling CPU Consumption.	627
	Controlling Parallel Execution of SQL Statements	630
	Setting a Memory Limit for SQL Statements.	633
	Managing Peak Load (Admission Control).	636
	Managing Workload with Workload Classes.	640
	Example Workload Management Scenarios.	655
6.10	Scheduling of Recurring Administration Tasks.	658
6.11	Getting Support.	659
	Diagnosis Files.	661
	Traces.	665
	Troubleshooting an Inaccessible or Unresponsive SAP HANA System	684
	Collecting Diagnosis Information for SAP Support.	687
	Problem Analysis Using hdbcons.	698
	Open a Support Connection.	698
7	Security Administration and User Management.	700
7.1	Monitoring Critical Security Settings in SAP HANA Cockpit.	700
	View Status of Security Settings.	701
	Security Tiles and Links.	701
	Network Security Details.	704
7.2	Managing SAP HANA Users.	706
	Database Users.	707
	Operating System User <sid>adm.	714
	User Authentication and Single-Sign On.	715
	User Authorization.	738
	Provisioning Users.	775
7.3	Auditing Activity in the SAP HANA Database.	826
	Managing Auditing in the SAP HANA Cockpit.	827
	Managing Auditing in the SAP HANA Studio.	837
	Audit Trail Targets.	843
	Best Practices and Recommendations for Creating Audit Policies.	845

7.4	Managing Data Encryption in SAP HANA.	847
	Server-Side Data Encryption Services.	848
	SAP HANA Client Secure User Store (hdbuserstore).	879
	Client-Side Data Encryption.	880
7.5	Managing Client Certificates.	900
	In-Database Certificate Management Workflow.	903
	Client Certificates.	904
	Certificate Collections.	904
	View Certificates in the Certificate Store.	906
	View Certificate Collections.	907
	Import a Trusted Certificate into the Certificate Store.	909
	Create a Certificate Collection.	910
	Set the Purpose of a Certificate Collection.	912
	Export a Client Certificate.	913
	SQL Statements and Authorization for In-Database Certificate Management (Reference).	913
7.6	Data Anonymization.	916
	Show Anonymization Views.	917
8	SAP HANA Lifecycle Management.	918
8.1	SAP HANA Platform Lifecycle Management.	919
	About the SAP HANA Database Lifecycle Manager (HDBLCM).	921
8.2	SAP HANA Application Lifecycle Management.	951
	Installing and Updating SAP HANA Products and Software Components in SAP HANA XS Classic Model.	952
	Installing and Updating Products and Software Components in SAP HANA XS Advanced Model	960
	Configuring SAP HANA Applications with the Process Engine.	989
8.3	SAP HANA Content.	994
	SAP HANA Archive Types.	994
	Deploy a Product Archive (*.ZIP).	995
	Deploy a Delivery Unit Archive (*.tgz).	996
9	Landscape Management and Network Administration.	997
9.1	Landscape Management.	997
	Copying and Moving a System Using Platform LCM Tools.	999
	Copying and Moving Tenant Databases Between Systems.	1004
	Copying a System using System Replication	1031
	Renaming a System.	1032
9.2	Network Administration.	1040
	Network Zones.	1041
	Ports and Connections.	1043
	Host Name Resolution.	1068

10	Availability and Scalability	1080
10.1	High Availability for SAP HANA	1080
	SAP HANA High Availability Support	1082
	Configuring SAP HANA System Replication	1089
	Setting Up Host Auto-Failover	1208
	Implementing a HA/DR Provider	1215
10.2	SAP HANA Database Backup and Recovery	1229
	Savepoints and Redo Logs	1230
	Points to Note About Backup and Recovery	1231
	Authorizations for Backup and Recovery	1244
	SAP HANA Backup	1245
	SAP HANA Recovery	1331
	Copying a Database Using Backup and Recovery	1374
	Planning Your Backup and Recovery Strategy	1393
	Reference: Backup Console (SAP HANA Studio)	1396
	Reference: Backup Alerts	1400
10.3	Scaling SAP HANA	1404
	Aspects of Scalability	1404
	Adding and Removing Hosts	1406
	Configuring Host Roles	1428
	Configuring the Network for Multiple Hosts	1438
	Scaling SAP HANA Extended Application Services, Classic Model	1448
	Starting and Stopping Distributed SAP HANA Systems Using SAPControl	1448
11	SAP HANA Deployment Infrastructure	1450
11.1	HDI Administration in Context	1452
11.2	HDI Administrator Roles	1454
11.3	The SQL API for SAP HANA Deployment Infrastructure (HDI)	1456
11.4	Enabling HDI in the Database	1460
	Enable HDI for a Specific Tenant on a Multi-Tenant Database	1461
	Create an HDI Administrator	1463
	Revoke the HDI Administrator Privileges	1464
11.5	Maintaining the HDI	1465
	HDI Container Group Administration	1466
	HDI Container Administration	1467
	HDI Container Schemas	1468
	Configure HDI Parameters	1469
	List Plug-in Libraries That Can Be Configured for a Container	1489
	Create a Container Group	1490
	Drop a Container Group	1491
	Grant Container Group Administrator Privileges to Another User	1491
	Revoke Container Group Administrator Privileges from a Container Group Administrator	1493

	Move a Container to Another Container Group.	1494
11.6	Maintaining HDI Container Groups.	1494
	Grant Container-Group Administrator Privileges to a User.	1496
	Revoke Container-Group Administrator Privileges from an Administrator User.	1497
	Create a Container.	1498
	Drop a Container.	1499
	Grant Container Administrator Privileges to a User.	1500
	Revoke Container Administrator Privileges from a User.	1501
	Grant a Support User Access to a Container.	1502
	Revoke Access to a Container from a Support User.	1504
	Export a Container for Copy Purposes.	1505
	Import a Container for Copy Purposes.	1506
11.7	Maintaining HDI Containers.	1508
	Grant HDI Container Administrator Privileges to a User.	1509
	Revoke HDI Container Administrator Privileges from a User.	1511
	Grant Access to the HDI Container Content-Development API.	1512
	Revoke Access to the HDI Container Content-Development API.	1513
	Grant Access to an HDI Container's Schema.	1515
	Revoke Access to an HDI Container's Schema.	1516
	Grant a User a Role from the Container's Schema.	1517
	Revoke a Role from the Container's Schema.	1518
	List All Currently Configured Build Plug-in Libraries Available to a Container.	1519
	Configure the Default Build Plug-in Libraries Available to a Container.	1520
	Configure a Custom Set of Build Plug-in Libraries Available to a Container.	1521
	Configure Container Parameters.	1523
	Cancel a Running Make Operation in a Container.	1524
12	Application Run-Time Services.	1526
12.1	Maintaining the SAP HANA XS Classic Model Run Time.	1526
	SAP HANA XS Classic Administration Tools.	1527
	SAP HANA XS Classic Administration Roles.	1529
	SAP HANA XS Classic Configuration Parameters.	1532
	Maintaining Application Runtime Configurations.	1538
	Managing Trust Relationships.	1552
	Maintaining SAML Providers.	1559
	Maintaining SMTP Server Configurations.	1568
	Maintaining HTTP Access to SAP HANA.	1574
	Maintaining Single Sign-On for SAP HANA XS Applications.	1583
	Maintaining User Self Service Tools.	1593
	Scheduling XS Jobs.	1618
	Maintaining Translation Text Strings.	1632
	Maintaining HTTP Traces for SAP HANA XS Applications.	1639

12.2	Maintaining the SAP HANA XS Advanced Model Run Time.	1647
	XS Advanced Platform Components.	1648
	Maintaining the XS Advanced Run-time Environment with a Command-Line Interface.	1652
	Maintaining the XS Advanced Run-time Environment with a Graphical User Interface.	1753
	XS Advanced User Management.	1814
	XS Advanced System Configuration Parameters.	1824
	Backup and Recovery in XS Advanced.	1831
	Logging and Auditing in XS Advanced.	1836
	Platform Sizing in XS Advanced.	1839
	Configuring the XS Advanced Platform Router.	1842
	Maintaining Single Sign-On for XS Advanced Applications.	1845
13	Data Access.	1848
13.1	SAP HANA Smart Data Access.	1849
	Setting Up ODBC Drivers.	1850
	Remote Source Credential Types.	1860
	Linked Database Overview.	1861
	Privilege Maintenance.	1862
	Creating a Remote Source	1863
	Modifying a Remote Source.	1909
	Dropping a Remote Source.	1910
	Listing Remote Sources.	1911
	Managing Secondary Credentials.	1912
	Managing Single Sign-On (SSO) with Kerberos.	1915
	Enabling Read-Write Access to a Remote Source.	1919
	Managing Virtual Tables.	1920
	Managing Linked Database.	1929
	EXPORT/IMPORT Virtual Tables.	1931
	Monitor Remote Connections and Statements.	1934
	Data Type Support.	1941
	Functions Pushed Down to Remote Sources.	1941
	Synonyms.	1942
	Statistics on Virtual Tables and Linked Database.	1943
	Pool of Remote Connections	1955
	Results Caching for Virtual Tables and Linked Database.	1957
	SAP HANA Automatic Failover Support.	1958
	Safe Mode for ODBC Connections.	1959
	Setting Session Specific Information for Connections.	1960
	Smart Data Access System Parameters.	1960
	Troubleshooting Smart Data Access.	1962
13.2	SAP HANA Hadoop Integration.	1962

14	SAP HANA HDBSQL (Command-Line Reference)	1963
14.1	SAP HANA HDBSQL Options	1963
14.2	Log On to a Database	1971
14.3	Run Commands	1973
14.4	Run Long Commands in Multiple-Line Mode	1975
14.5	Edit Long Commands in an External File	1976
14.6	Redirect Results to a File	1977
14.7	Substitution Variables	1978

1 SAP HANA Administration Guide

The SAP HANA Administration Guide is the central operations documentation for the on-premise deployment of the SAP HANA platform.

i Note

The *SAP HANA Administration Guide* does not cover administration tasks related to some additional capabilities that may be installed in the SAP HANA system such as SAP HANA dynamic tiering and SAP HANA streaming analytics. For more information about the administration of these capabilities, see the relevant documentation on SAP Help Portal.

2 Administration Information Map

In addition to the SAP HANA Administration Guide, several other documents in the SAP HANA platform documentation set provide administrators with important information.

SAP HANA Administration Guide

This guide is the entry point for all information related to the ongoing operation and maintenance of an on-premise deployment of the SAP HANA platform. It contains information about various administration tasks related to the following main areas:

- Administration and monitoring at the landscape, system, and database level
- Monitoring and configuration of security-related settings
- Landscape management and network administration
- Administration of the SAP HANA XS runtime environment
- Setup and management of the SAP HANA Deployment Infrastructure (HDI)
- High availability (including backup and recovery) and scalability
- Data access and integration with SAP HANA data provisioning tools and technologies

The *SAP HANA Administration Guide* also includes information on using the following native SAP HANA administration tools:

- SAP HANA cockpit

→ Tip

For the documentation of the latest SAP HANA cockpit support package (SP), see https://help.sap.com/viewer/p/SAP_HANA_COCKPIT

- SAP HANA studio
- SAP HANA database lifecycle manager (HDBLCM)
- SAP HANA `hdsq1` command line
- SAP HANA XS administration tools
- SAP HANA application lifecycle management

SAP HANA Troubleshooting and Performance Analysis Guide

This guide describes what steps you can take to identify and resolve specific performance issues and what you can do to enhance the performance of your SAP HANA database in areas such as:

- Host resources (CPU, memory, disk)
- Size and growth of data structures
- Transactional problems
- SQL statement performance
- Security, authorization, and licensing
- Configuration

Open the [SAP HANA Troubleshooting and Performance Analysis Guide](#)

SAP HANA Tenant Databases Operations Guide

This guide brings together all the information required for the operation of an SAP HANA multitenant system, including:

- Overview of architecture and concepts of multitenant systems
- Creating and configuring tenant databases
- Monitoring and managing tenant databases
- Copying and moving tenant databases

Open the [SAP HANA Tenant Databases Operations Guide](#)

SAP HANA Master Guide

This guide is the entry point for planning the installation of your SAP HANA system landscape. It provides you with overview information on the aspects such as:

- Use cases and scenarios that SAP HANA can be used in from an application point of view
- Deployment options for SAP HANA on-premise or in the cloud
- Implementation and operation activities during the lifecycle of SAP HANA

Open the [SAP HANA Master Guide](#)

SAP HANA SQL and System Views Reference

The *SAP HANA SQL Reference* describes all SQL data types, predicates, operators, expressions, functions, statements, and error codes. The *SAP HANA System Views Reference* describes all system views. You can use the information in this guide to perform the following typical tasks:

- Monitor the current status of the SAP HANA system and database by querying monitoring views
- Analyze and diagnose historical monitoring data by querying statistics views
- Configure the database using SQL commands

Open the [SAP HANA SQL and System Views Reference](#)

Administration Guides for Additional SAP HANA Capabilities

Separate administration information is available for the following additional capabilities that may be installed in your SAP HANA system:

- SAP HANA accelerator for SAP ASE
- SAP HANA data warehousing foundation
- SAP HANA dynamic tiering
- SAP HANA remote data sync
- SAP HANA smart data integration and SAP HANA smart data quality
- SAP HANA streaming analytics
- SAP HANA real-time replication with SAP Landscape Transformation Replication Server

For more information, see the relevant documentation on SAP Help Portal.

i Note

The topics listed above for each guide are not intended to be exhaustive but representative.

Target Audiences

Document	Target Audience	Content Type
SAP HANA Administration Guide	Technology consultants, system administrators	Task- and role-oriented
SAP HANA Tenant Databases Operations Guide	Technology consultants, system administrators	Concept, task- and role-oriented
SAP HANA Troubleshooting and Performance Analysis Guide	Technology consultants, system administrators	Troubleshooting, root-cause analysis
SAP HANA Master Guide	Technology consultants, security consultants, system administrators	Concept and overview
SAP HANA SQL and System Views Reference	Technology consultants, security consultants, system administrators	Reference

Additional Documentation Resources

Product Documentation

For more information about the SAP HANA landscape, including installation and security, see https://help.sap.com/viewer/p/SAP_HANA_PLATFORM.

SAP Notes

SAP Note	Title
2380229	SAP HANA Platform 2.0 – Central Note
1730928	Using external software in an SAP HANA appliance
1730929	Using external tools in an SAP HANA appliance
1730930	Using anti-virus software in an SAP HANA appliance
1730996	Non-recommended external software and software versions
1730997	Non-recommended versions of anti-virus software
1730998	Non-recommended versions of backup tools
1730999	Configuration changes in SAP HANA appliance
1731000	Non-recommended configuration changes

Other Information

For more information about specific topics, see the quick links in the table below.

Content	SAP Service Marketplace or SDN Quick Link
Related SAP Notes	https://support.sap.com/notes
Released platforms	https://apps.support.sap.com/sap/support/pam

Content	SAP Service Marketplace or SDN Quick Link
SAP Solution Manager community	https://go.sap.com/community/topic/solution-manager.html 
SAP NetWeaver community	https://go.sap.com/community/topic/netweaver.html 
In-memory computing community	https://go.sap.com/community/topic/hana.html 

3 Database Administration Tasks at a Glance

Overview of key tasks for the ongoing operation and maintenance of the SAP HANA database

Initial Administration Tasks

After the initial setup and initial data load, it is strongly recommended that you perform a full data and file-system backup (including configuration backup). For more information, see the section on database backup and recovery.

i Note

In replication scenarios with SAP Landscape Transformation Replication Server, do not switch off log writing during the initial data load from SAP ERP into the SAP HANA database. There is no system table or log file that records the information that log writing has been switched off, so it is not possible to check whether log writing has been switched on or off.

Regular Administration Tasks

- Monitor the health of your SAP HANA system using, for example, the SAP HANA cockpit. The most important system information to review is:
 - Overall system status
 - Status of database services, for example, name server and index server
 - General system information (software versions and so on)
 - Alerts generated by the statistics service
 - Usage of important system resources: memory, CPU and diskFor more information, see *Monitoring the SAP HANA Database*.
- Perform regular data backups, including configuration backups. There are no specific guidelines for backup frequency, which depends on the usage scenario, but for general guidelines, see *Planning Your Backup and Recovery Strategy*.
- Avoid the log backup area becoming full by archiving old log backups to a different location.

⚠ Caution

Do not delete log segments at the operating system level, as the log area will become unusable and the database may stop working.

For more information, see *Housekeeping for Backup Catalog and Backup Storage*.

- Monitor disk space used for diagnosis files and delete files that are no longer needed.

On-Demand Administration Tasks

- In the event of problems with the SAP HANA database, you can check log and trace files for errors. You can also activate and configure several trace types.
For more information, see *Diagnosis Files* and *Configure Traces*.
- Before updating SAP HANA, perform a data backup including configuration files. This allows for the recovery of the system in the event the software update fails.

Related Information

[SAP HANA Database Backup and Recovery \[page 1229\]](#)

[Monitoring the SAP HANA Database \[page 317\]](#)

[Housekeeping for Backup Catalog and Backup Storage \[page 1259\]](#)

[Planning Your Backup and Recovery Strategy \[page 1393\]](#)

[Diagnosis Files \[page 661\]](#)

[Configure Traces in SAP HANA Studio \[page 683\]](#)

[Configure Tracing in the SAP HANA Database Explorer \[page 682\]](#)

4 SAP HANA System Architecture Overview

An SAP HANA system comprises multiple isolated databases and may consist of one host or a cluster of several hosts.

An **SAP HANA system** is identified by a single system ID (SID) and contains one or more tenant databases and one system database. Databases are identified by a SID and a database name. From the administration perspective, there is a distinction between tasks performed at system level and those performed at database level. Database clients, such as the SAP HANA cockpit, connect to specific databases.

The **SAP HANA XS advanced application server** is a layer on top of SAP HANA that provides the platform for running SAP HANA-based Web applications. It is an integral part of the SAP HANA system.

A system may consist of one host or a cluster of several hosts. This is referred to as a **multiple-host, distributed system, or scale-out system** and supports scalability and availability.

The following sections provide overview information about these aspects of system architecture.

Related Information

[Server Architecture of Tenant Databases \[page 17\]](#)

[Server Architecture of SAP HANA XS Advanced Runtime Platform \[page 22\]](#)

[Multiple-Host \(Distributed\) Systems \[page 25\]](#)

[SAP HANA Services \[page 30\]](#)

4.1 Server Architecture of Tenant Databases

An SAP HANA database consists of multiple servers, for example, name server, index server, preprocessor server, and so on. The databases in an SAP HANA system run different combinations of these servers. The most important server is the index server. It contains the actual data stores and the engines for processing the data and runs in every tenant database.

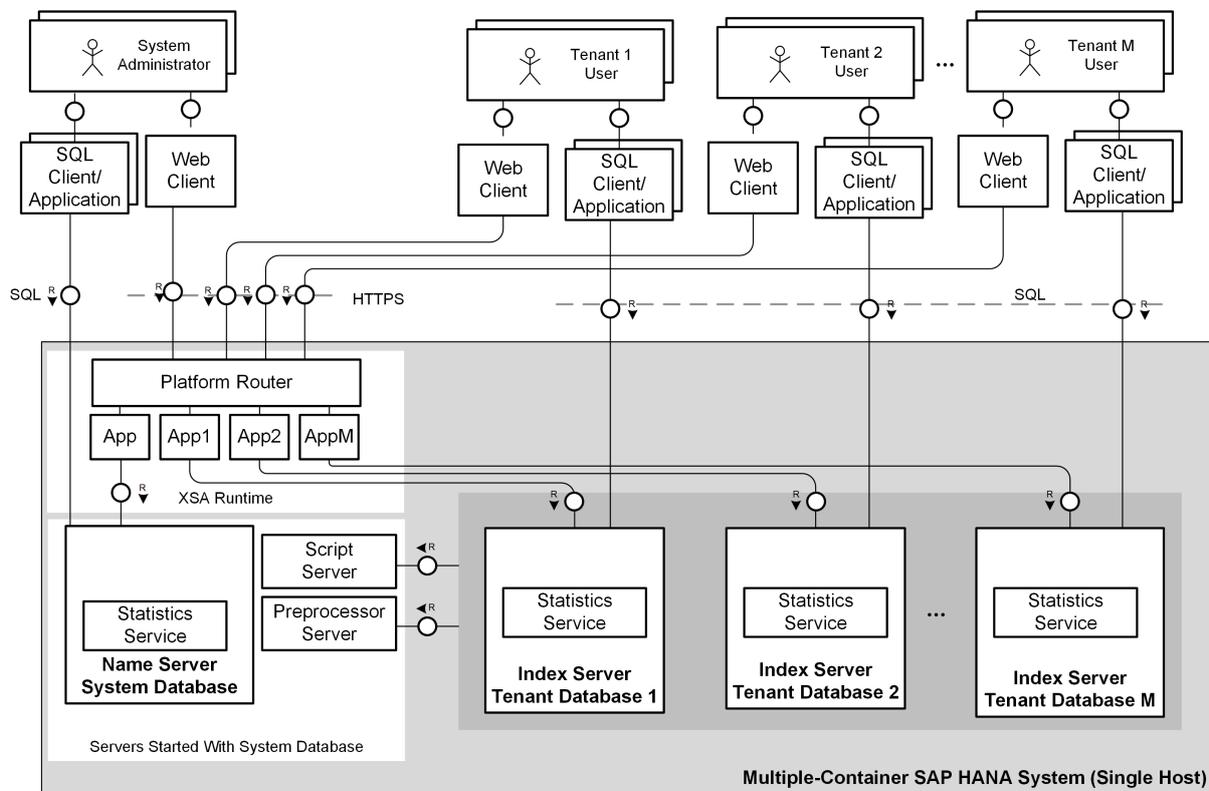
Only the **system database** runs the name server. The name server contains landscape information about the system as a whole, including which tenant databases exist. It also provides indexserver functionality for the system database. The name server of the system database in a tenant database system does not own topology information, that is, information about the location of tables and table partitions in databases. Database-related topology information is stored in the relevant tenant database catalog.

Tenant databases require only an own index server. Servers that do not persist data, such as the compile server and the preprocessor server, run on the system database and serve all databases.

i Note

For a full list and description of all SAP HANA servers, see *SAP HANA Services*.

The following figure shows a sample system with three databases (system database and three tenant databases) on a single host.



Single-Host SAP HANA System with Tenant Databases

Note

If the SAP HANA XS classic server is available, it runs embedded in the (master) index server of the tenant database by default, although it can be added as a separate service if necessary. The **SAP Web Dispatcher**, which runs as a separate database service on the host of the system database, is used to route incoming HTTP requests from clients to the correct XS classic server based on virtual host names. This is part of network configuration. In addition to the system-internal Web Dispatcher, you can implement an external Web Dispatcher for load distribution. See the section on using the SAP Web Dispatcher for load balancing with tenant databases.

Related Information

[SAP HANA Services \[page 30\]](#)

[Connections from Database Clients and Web Clients to SAP HANA \[page 1043\]](#)

[Port Assignment in Tenant Databases \[page 1050\]](#)

[Scale-Out Architecture of Tenant Databases \[page 26\]](#)

[Using SAP Web Dispatcher for Load Balancing with Tenant Databases \[page 280\]](#)

4.1.1 Tenant Databases

SAP HANA supports multiple isolated databases in a single SAP HANA system. These are referred to as tenant databases.

An SAP HANA system is capable of containing more than one tenant database.

A system always has exactly one system database, used for central system administration, and any number of tenant databases (including zero). An SAP HANA system is identified by a single system ID (SID). Databases are identified by a SID and a database name. From the administration perspective, there is a distinction between tasks performed at system level and those performed at database level. Database clients, such as the SAP HANA cockpit, connect to specific databases.

All the databases share the same installation of database system software, the same computing resources, and the same system administration. However, each database is self-contained and fully isolated with its own:

- Set of database users
- Database catalog
- Repository
- Persistence
- Backups
- Traces and logs

Although database objects such as schemas, tables, views, procedures, and so on are local to the database, cross-database SELECT queries are possible. This supports cross-application reporting, for example.

Related Information

[Server Architecture of Tenant Databases \[page 17\]](#)

[Scale-Out Architecture of Tenant Databases \[page 26\]](#)

[The System Database \[page 19\]](#)

[Administration of Tenant Databases \[page 20\]](#)

4.1.2 The System Database

The system database is created during either installation or conversion from a single-container system to a tenant database system. The system database contains information about the system as a whole, as well as all its tenant databases. It is used for central system administration.

A system has exactly one system database. It contains the data and users for system administration. System administration tools, such as the SAP HANA cockpit, can connect to this database. The system database stores overall system landscape information, including knowledge of the tenant databases that exist in the system. However, it doesn't own database-related topology information, that is, information about the location of tables and table partitions in databases. Database-related topology information is stored in the relevant tenant database catalog.

Administration tasks performed in the system database apply to the system as a whole and all of its databases (for example, system-level configuration settings), or can target specific tenant databases (for example, backup of a tenant database). For more information, see *Administration of Tenant Databases*.

Things to Remember About the System Database

- The system database does not have the same functionality as a tenant database.
- The system database is not a database with full SQL support.
- The system database cannot be distributed across multiple hosts, in other words, scale-out is not possible.
- If you need a full-featured SAP HANA database, you always have to create at least one tenant database.
- The system database does not support Application Function Libraries (AFL) and SAP liveCache applications.
- Cross-database access between the system database and a tenant database is not possible. The system database can show monitoring data from tenant databases (views in the schema SYS_DATABASES) but can never show actual content from tenant databases.
- The system database cannot be copied or moved to another host.
- SAP HANA options can only run in tenant databases.
- Tenant-specific configurations cannot be set in the system database. Only global settings are allowed.
- Features can only be restricted or disabled at high level for tenant databases.

Related Information

[Administration of Tenant Databases \[page 20\]](#)

[Configuring Memory and CPU Usage for Tenant Databases \[page 266\]](#)

4.1.3 Administration of Tenant Databases

In SAP HANA systems there is a distinction between administration tasks performed at system level and those performed at database level.

System Versus Database Administration

Tenant database systems have two levels of administration.

Some administration tasks are performed in the system database and apply globally to the system and all its databases. They include for example:

- Starting and stopping the whole system
- Monitoring the system

- Configuring parameters in configuration (*.ini) files at system level
- Setting up and configuring tenant databases, for example:
 - Creating and dropping tenant databases
 - Disabling features on tenant databases
 - Configuring system- and database-specific parameters in configuration (*.ini) files
 - Scaling out tenant databases by adding services
- Backing up tenant databases
- Recovering tenant databases

Some administration tasks are performed in the tenant database and apply only to that database. They include for example:

- Monitoring the database
- Provisioning database users
- Creating and deleting schemas, tables, and indexes in the database
- Backing up the database
- Configuring database-specific parameters in configuration (*.ini) files

Administration Tools

Several tools are available for the administration of SAP HANA. While all tools support database-level administration, system-level administration of tenant databases requires the SAP HANA cockpit (for example, monitoring availability of tenant databases, creating and deleting tenant databases).

For more information about the SAP HANA cockpit and other administration tools, see the section on administration tools in the *SAP HANA Administration Guide*.

Related Information

[Tenant Databases \[page 19\]](#)

[The System Database \[page 19\]](#)

[Creating and Configuring Tenant Databases \[page 189\]](#)

[SAP HANA Administration Tools \[page 39\]](#)

[Monitoring and Managing Tenant Databases \[page 243\]](#)

4.2 Server Architecture of SAP HANA XS Advanced Runtime Platform

SAP HANA extended application services, advanced model (XS advanced for short) provides a comprehensive platform for the development and execution of micro-service oriented applications, taking advantage of SAP HANA's in-memory architecture and parallel execution capabilities.

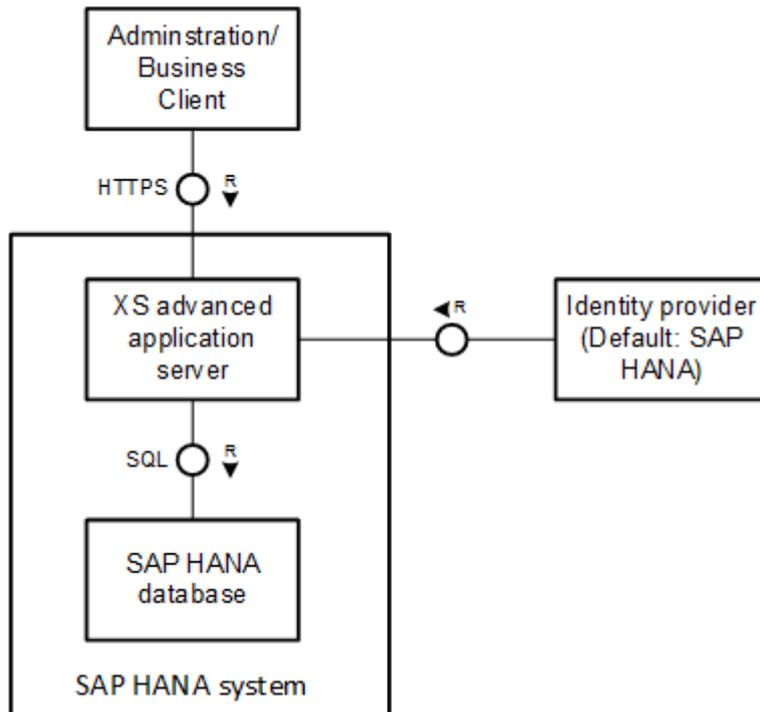
About SAP HANA XS Advanced

SAP HANA XS advanced offers a rich set of embedded services that enable an end-to-end support for web-based applications including lightweight web servers, persistency services, and a configurable identity provider. Furthermore, the platform supports polyglot application development with a core set of pre-deployed runtimes that are accepted as industry standard, for example, node.js or JavaEE.

Although the built-in runtimes come with first-class development and monitoring support, the platform has an open architecture that allows you to add custom runtimes. This high flexibility makes it essential that you put a strong focus on security concepts, not only when configuring and setting up the infrastructure, but also throughout operating the system.

Architecture Overview

As illustrated in the following diagram, the basic system architecture has a classic 3-tier approach:

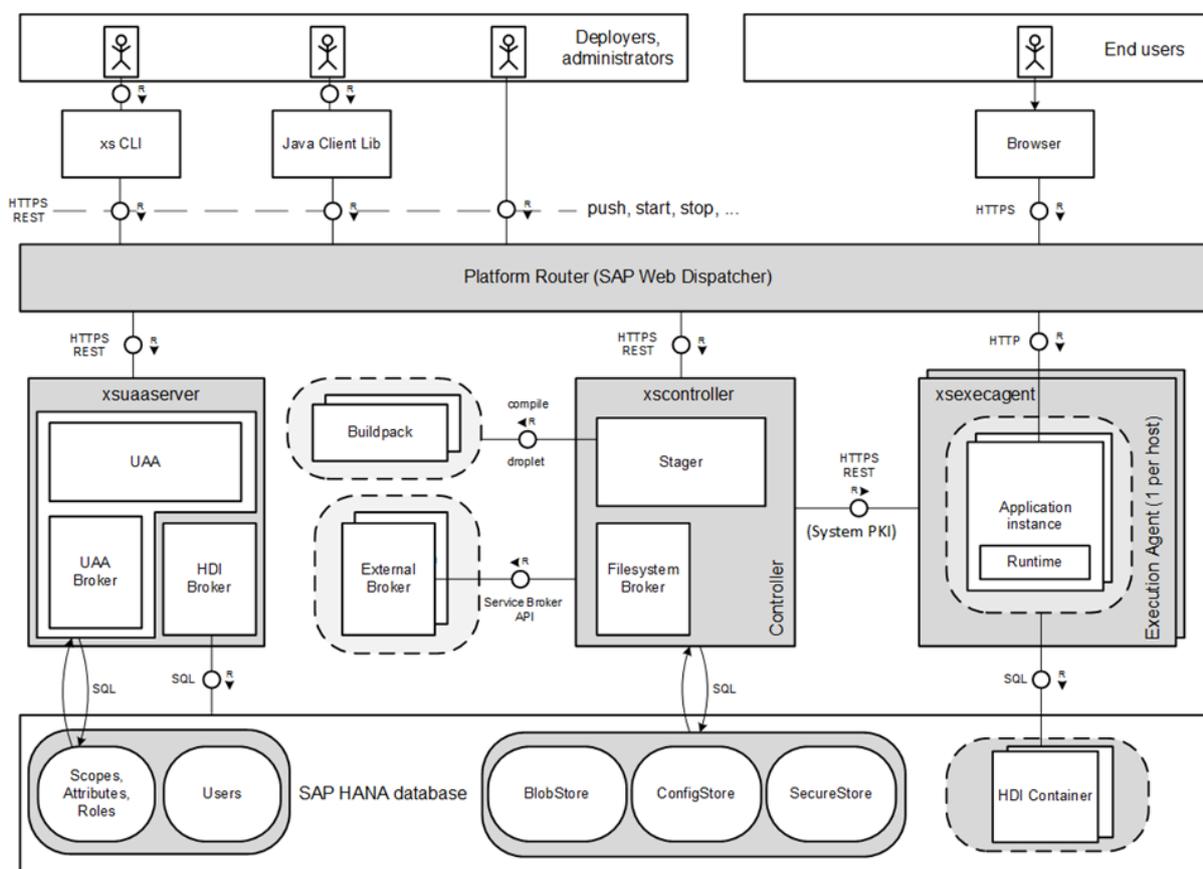


3-Tier Architecture of SAP HANA with XS Advanced

There is a distinction between the overall SAP HANA system and the SAP HANA XS advanced application server. The SAP HANA system refers to the entire SAP HANA platform part of the integrated solution. The SAP HANA XS advanced application server describes only the runtime platform as an integral part of the solution. All services of the SAP HANA system share the same system identifiers (that is, instance number and SID) and are controlled by the `hdbdaemon` service.

The third tier, represented by an SAP HANA database, provides persistency services, that is, data storage. In contrast, the application server components in the middle tier are responsible for deploying, running, and monitoring the applications. Most security-related features such as authentication, authorization, and auditing are primarily enforced in this layer. End users interact on the client layer with system or business users that are authenticated by an identity provider (IdP), which is SAP HANA user management by default. However, both the server components and the applications themselves access the SAP HANA database only through technical database users that the platform generates implicitly. Direct access to the database is only intended for database administration purposes.

The following diagram provides a more detailed overview of the technical system landscape of the XS advanced application server. All relevant components and storages used by the application server layer are highlighted with a gray background.



Technical System Landscape of XS Advanced Application Server

The XS advanced application server relies on the following SAP HANA services contributing to the integrated platform solution:

1. `xscontroller` (Controller, FileSystem Broker, Platform Router)
2. `xsexecagent` (Execution Agent)
3. `xsuaaserver` (UAA, UAA Broker and HDI Broker)

Administration of the XS Advanced Runtime

A number of administration tools are available to enable you to maintain and manage the various components of the XS advanced runtime environment. For more information, see the section on maintaining the SAP HANA XS advanced model run time.

Related Information

[SAP HANA Services \[page 30\]](#)

[Connections for SAP HANA Extended Application Services, Advanced Model \[page 1060\]](#)

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1647\]](#)

4.3 Multiple-Host (Distributed) Systems

An SAP HANA system can be distributed across multiple hosts for reasons of scalability and availability.

A multiple-host or distributed SAP HANA system is a system that is installed on more than one host. Otherwise, it is a single-host system.

An SAP HANA system installed on multiple hosts is identified by a single system ID (SID). It is perceived as one unit from the perspective of the administrator, who can install, update, start up, or shut down the system as a whole. The different databases of the system share the same metadata and requests from client applications can be transparently dispatched.

The main reason for distributing a system across multiple hosts is **scale-out**. A multiple-host system can overcome hardware limitations of a single physical server, and it can distribute the load between multiple servers. Distributing a system also supports **failover**. One or more hosts can be configured to work in standby mode, so that if an active hosts fails, a standby host automatically takes its place. The index servers on standby hosts do not contain any data and do not receive any requests.

For more information about hosts, including host roles, fail-over configuration, and storage options, see the *SAP HANA Server Installation and Update Guide*.

Distributing Data

In a multiple-host system each index server is usually assigned to its own host for maximum performance. SAP HANA supports different ways of distributing data across the hosts:

- Different tables can be assigned to different index servers.
- A table can be split, or partitioned, in a way that different rows of the table are stored on different index servers
- A table can be replicated to multiple index servers, for better query and join performance.

When you create new tables or partitions, data is distributed to the available hosts by the system. By default a 'round-robin' distribution method is used, but tables can also be positioned by using table placement rules or by specifying a host and port number with the SQL CREATE TABLE statement in the location clause; this gives complete control over the positioning of individual tables.

Specific applications may have predefined table distribution rules and in some cases configuration files and documentation are available in SAP Notes to help you to set up the necessary partitioning and table placement rules.

For more information, see the sections on table placement, table partitioning, and table replication.

Distributed Execution

Database clients may send their requests to any index server on any host in a distributed system. If the contacted index server does not own all of the data involved, it delegates the execution of some operations to other index servers, collects the result, and returns it to the database client. The SAP HANA client library

supports load balancing and minimizes communication overhead by selecting connections based on load data and routing statements based on information about the location of data.

For more information, see the sections on connecting to SAP HANA databases and servers and statement routing in the *SAP HANA Client Interface Programming Reference*.

Related Information

[Multiple-Host System Concepts \[page 1406\]](#)

[Connections for Distributed SAP HANA Systems \[page 1053\]](#)

[Scale-Out Architecture of Tenant Databases \[page 26\]](#)

[Scale-Out Architecture of SAP HANA XS Advanced Runtime Platform \[page 28\]](#)

[High Availability for SAP HANA \[page 1080\]](#)

[Scaling SAP HANA \[page 1404\]](#)

[Table Placement \[page 574\]](#)

[Table Partitioning \[page 542\]](#)

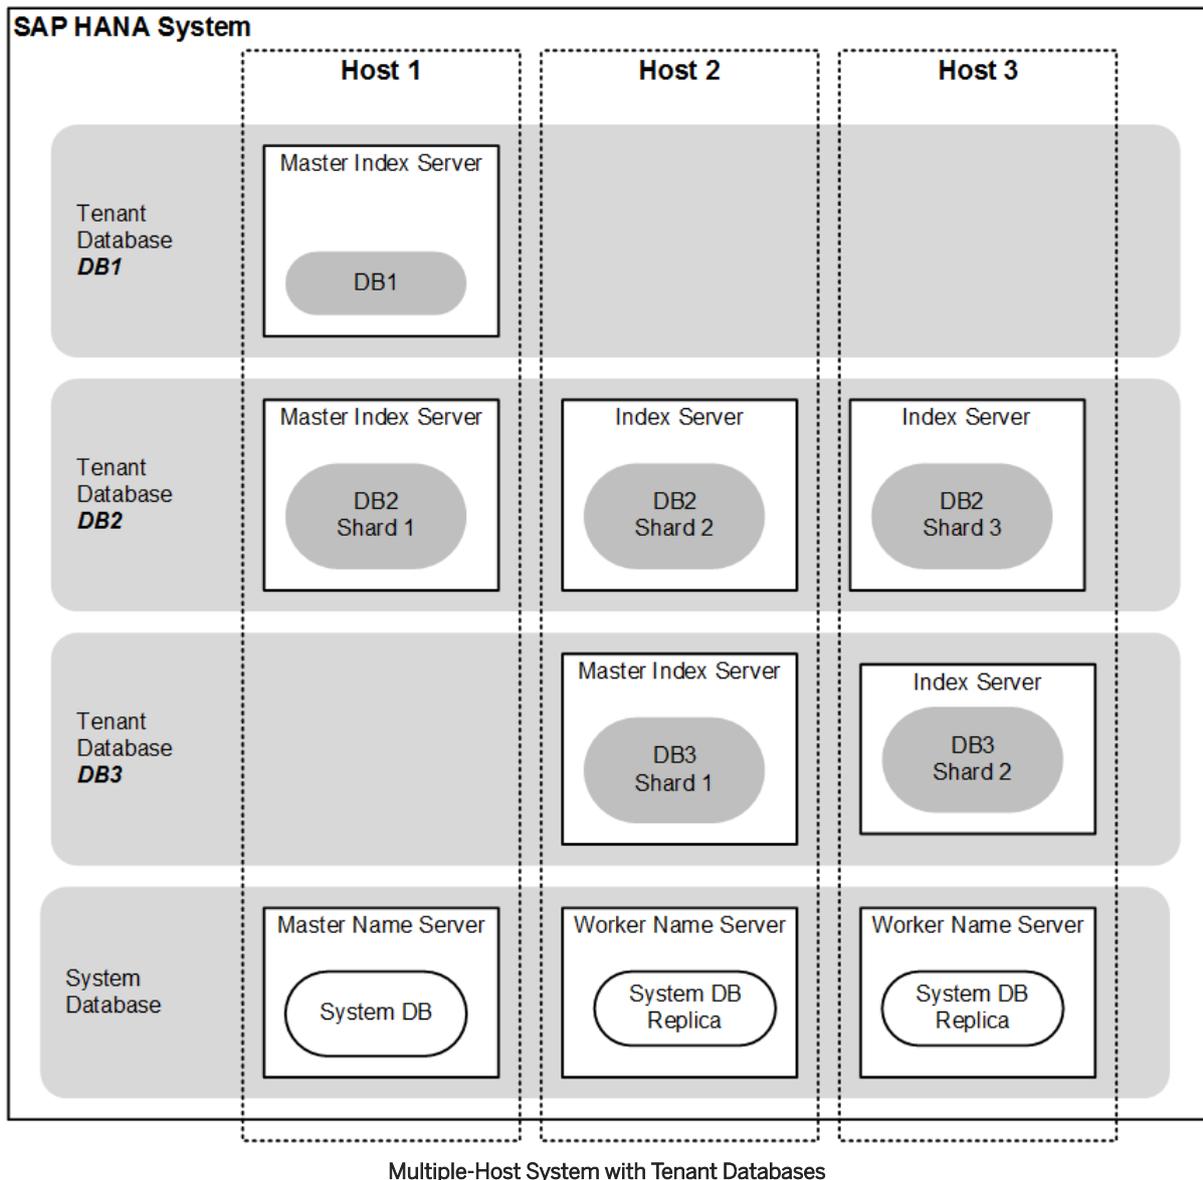
[Table Replication \[page 582\]](#)

4.3.1 Scale-Out Architecture of Tenant Databases

Tenant databases can be distributed across several hosts in a multiple-host system.

To ensure system availability, an instance of the system database runs on all hosts (worker and standby) in a single master and multiple workers configuration. Tenant databases can be created on worker hosts and existing databases can be scaled out through the addition of services. If a host fails, the standby instance will fail over all active databases and their services. Like in a single-host system, the master candidate for a failing host is determined. On that host the system database is restarted, if necessary. Up to three hosts can be configured to act as the master host of a system. These three hosts can be set up in the clients with the database name to be reconnected to a tenant database even in the case of a host auto-failover of the master host with the system database.

The following figure shows a tenant database system with three tenant databases distributed across three hosts. Tenant database DB1 has only one index server on host 1, while DB2 and DB3 are distributed across several hosts. Tenant database DB2, for example, is divided into three database shards, each of them with its own index server on a different host. In this context, a database shard is the union of all tables, partitions and replicas of one database that reside on one index server. Tenant database DB3 consists of two shards, one on host 2 and one on host 3. System administrators can specify the host when they create the tenant database, or they can let SAP HANA choose an appropriate host based on load-balancing algorithms.



Scale-Out Recommendations

When planning your SAP HANA deployment with tenant databases, various options exist with regard to scale-up versus scale-out.

In general, scaling up offers some performance advantages over scaling out, as memory access is local and minor overhead associated with inter-node network communication is avoided.

Note the following with regard to scale-out:

- It is possible to distribute tenant databases across several hosts in a scale-out system.
- The primary reason to distribute tenant databases generally is when their size is larger than the capacity of a single host. However, other reasons for distributing tenant database may exist, for example, a large SAP Business Warehouse (BW) system requires a scale-out configuration in accordance with its sizing rules.
- If tenant databases are distributed in a scale-out configuration due to sizing requirements, caution is advised when deploying additional tenant databases on the same host as a distributed tenant database

shard. The rationale is this: Workload in distributed scenarios can be somewhat volatile and less predictable. Therefore in many cases, it can be advantageous to dedicate maximum resources of the host to the distributed tenant database shard in order to maintain expected performance.

- In certain cases, more than one distributed tenant database shard may share the same host. In these cases, in order to dedicate maximum resources for a master node (for performance reasons), it is advisable to avoid deploying other tenant databases on the master node. For example, the following deployment should offer performance advantages:
 - Host 1: Master for tenant database 1
 - Host 2: Worker for tenant database 1 and worker for tenant database 2
 - Host 3: Master for tenant database 2
 - Host 4: Standby host for failover

Related Information

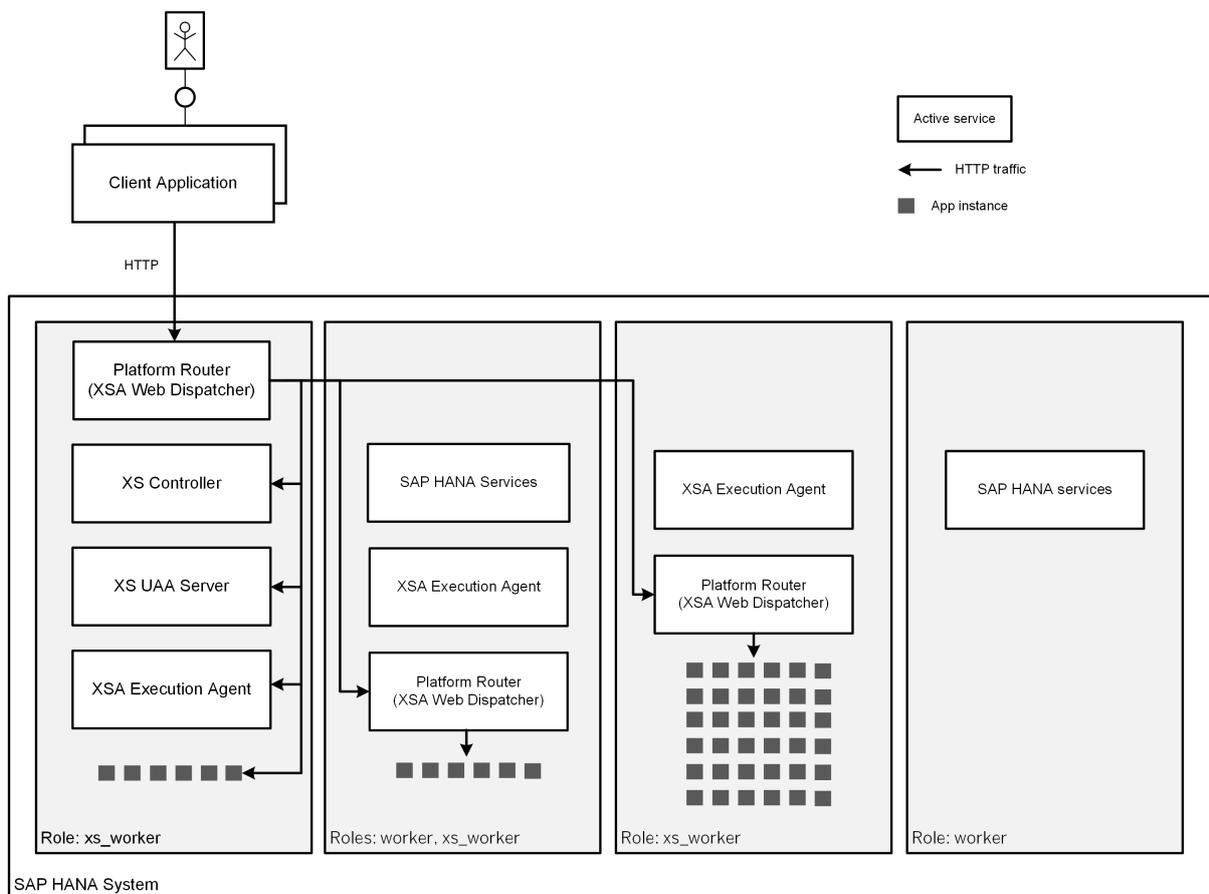
[Scaling SAP HANA \[page 1404\]](#)

4.3.2 Scale-Out Architecture of SAP HANA XS Advanced Runtime Platform

A multiple-host system that includes the SAP HANA XS advanced runtime can be flexibly configured to optimize load balancing and support failover.

During the installation of a multiple-host system with SAP HANA XS advanced, additional host roles are assigned for XS advanced. By default all worker hosts are configured to act as XS worker hosts; that is, they are additionally assigned the role `xs_worker`. However, it is also possible to configure dedicated hosts for XS advanced during installation.

The following figure shows a multiple-host system with two dedicated XS hosts and one shared host.



Multiple-Host System with Dedicated XS Worker Hosts and Shared Hosts

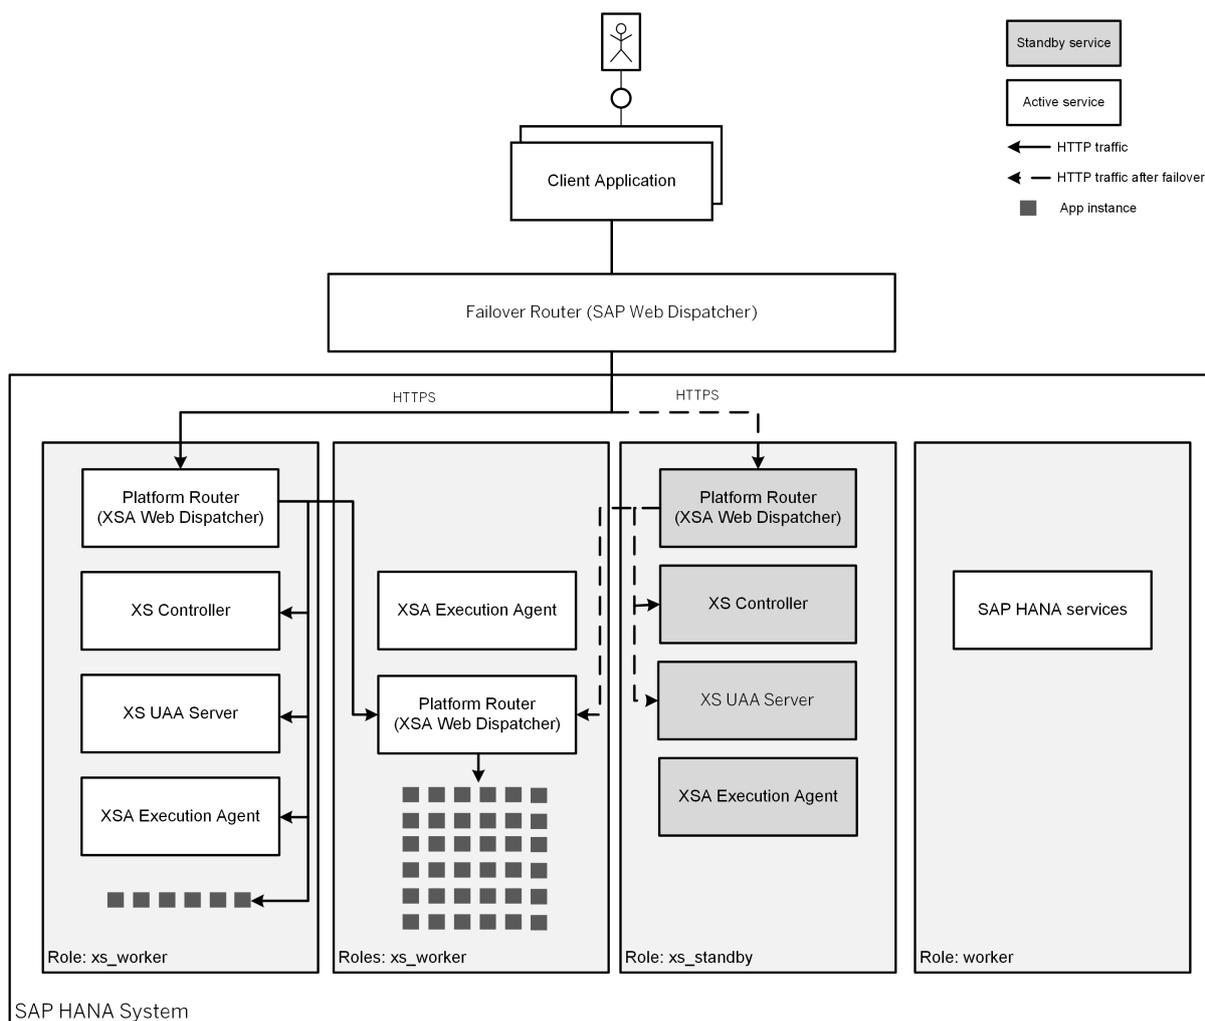
Additionally, the XS advanced runtime can be set up behind a reverse proxy (for example, a load balancer). For more information, see the section on installing the XS advanced runtime in the *SAP HANA Server Installation and Update Guide*.

Failover Configuration

SAP HANA XS advanced integrates into the standard failover mechanisms of SAP HANA: If a worker host fails, a standby host takes over. If a host is both a worker and an XS worker, then the regular standby host takes over. However, a dedicated XS worker host must have a dedicated XS standby host, that is a host with the role `xs_standby`.

In addition, to support failover an external router (for example, SAP Web Dispatcher) must be set up to route requests to the XS advanced hosts active after failover. For more information, see the section on high availability for the XS advanced runtime.

The following figure shows a multiple-host system with two dedicated XS worker hosts, one dedicated XS standby host and SAP Web Dispatcher as failover router.



Multiple-host System Configured for XS Advanced Failover

Related Information

[Host Auto-Failover Setup with XS Advanced Runtime \[page 1213\]](#)

[Connections for SAP HANA Extended Application Services, Advanced Model \[page 1060\]](#)

4.4 SAP HANA Services

Overview of the most important server components of the SAP HANA database and the corresponding OS processes and services

→ Tip

For more information about the ports used by the components listed here, see the section on ports and connections.

- [Core Services \[page 31\]](#)
- [Optional Services \[page 32\]](#)
- [SAP HANA XS Advanced Services \[page 34\]](#)

Core Services

Server Component	OS Process	Service Name	Description
Name server	<code>hdbnameserver</code>	<code>nameserver</code>	The name server, which runs in the system database only, owns the information about the topology of the SAP HANA system, including knowledge of the tenant databases that exist in the system.
Index server	<code>hdbindexserver</code>	<code>indexserver</code>	The index server, which runs in every tenant database (but not the system database), contains the actual data stores and the engines for processing the data.
Compile server	<code>hdbcompileserver</code>	<code>compileserver</code>	The compile server performs the compilation of stored procedures and programs, for example, SQLScript procedures. It runs on every host and does not persist data. It runs in the system database and serves all tenant databases.
Preprocessor server	<code>hdbpreprocessor</code>	<code>preprocessor</code>	The preprocessor server is used by the index server to analyze text data and extract the information on which the text search capabilities are based. It runs in the system database and serves all tenant databases.
SAP Web Dispatcher	<code>hdbwebdispatcher</code>	<code>webdispatcher</code>	The Web Dispatcher processes inbound HTTP and HTTPS connections to XS classic services.
SAP start service	<code>sapstartsrv</code>	<code>sapstartsrv</code>	The SAP start service is responsible for starting and stopping the other services in the correct order. It also performs other functions, such as monitoring their runtime state.

Optional Services

In addition to the main servers mentioned above, the following optional servers may also be running:

Server Component	OS Process	Service Name	Description
Script server	hdbscriptserver	scriptserver	<p>The script server is used to execute application function libraries written in C++.</p> <p>The script server is optional and must be added manually to the database that requires it. For more information, see the section on adding a service to a tenant database.</p>
Document store server	hdbdocstore	docstore	<p>This server is required for the document store repository. The document store allows native operations on JSON documents and joins with other column or row store tables.</p> <p>The document store is optional and must be added manually to the database that requires it. For more information, see the section on adding a service to a tenant database.</p>
XS advanced runtime	<ul style="list-style-type: none"> • hdbxscontroller • hdbxsexeagent • hdixsuaaserver 	<ul style="list-style-type: none"> • xscontroller • xsexeagent • hdixsuaaserver 	<p>SAP HANA includes a run-time environment for application development: SAP HANA extended application services, advanced model (XS advanced). The SAP HANA XS advanced model represents an evolution of the application server architecture within SAP HANA by building upon the strengths (and expanding the scope) of previous SAP HANA extended application services, classic model (XS classic).</p> <p>The SAP HANA XS advanced runtime consists of several processes for platform services and for executing applications. For more information about the individual services, see the table below.</p> <p>The SAP HANA XS advanced runtime runs either on dedicated hosts or together with other SAP HANA components on the same host.</p>
SAP HANA Deployment Infrastructure (HDI) server	hdbdiserver	diserver	<p>HDI handles the deployment of design-time artifacts into the SAP HANA database. If XS advanced is installed in the system, HDI is already enabled. Otherwise, you must enable it manually.</p> <p>For more information, see the section on enabling HDI in the database.</p>

Server Component	OS Process	Service Name	Description
XS classic server	hdbxsengine	xsengine	<p>SAP HANA Extended Application Services, classic model (XS, classic) is the application server for native SAP HANA-based web applications. It is installed with the SAP HANA database and allows developers to write and run SAP HANA-based applications without the need to run an additional application server. SAP HANA XS is also used to run web-based tools that come with SAP HANA, for instance for administration, lifecycle management and development.</p> <p>XS classic is the original implementation of SAP HANA XS.</p> <p>The XS classic server can run as a separate server process or embedded within the index server.</p> <div style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>SAP HANA XS, classic and the SAP HANA repository are deprecated as of SAP HANA 2.0 SPS 02. For more information, see SAP Note 2465027.</p> </div>
Extended store server	hdbesserver	esserver	<p>The extended store server is part of SAP HANA dynamic tiering. It provides a high-performance disk-based column store for very big data up to the petabyte range.</p> <p>For more information, see the documentation for SAP HANA dynamic tiering on SAP Help Portal.</p>
Data provisioning server	hdbdpserver	dpserver	<p>The data provisioning server is part of SAP HANA smart data integration. It provides capabilities such as data provisioning in real time and batch mode, real-time data transformations, data quality functions, adapters for various types of remote sources, and an adapter SDK for developing additional adapters.</p> <p>For more information, see the documentation for SAP HANA smart data integration on SAP HANA Portal.</p>
Streaming cluster	hdbstreamingserver	streamingserver	<p>The streaming cluster is part of SAP HANA streaming analytics. Streaming analytics extends SAP HANA with capabilities of SAP Event Stream Processor for consuming data streams and complex event processing.</p> <p>For more information, see the documentation for streaming analytics on SAP Help Portal.</p>

Server Component	OS Process	Service Name	Description
Accelerator for SAP ASE	hdbetsserver	etsserver	<p>The SAP ASE server is part of SAP HANA accelerator for SAP ASE. It provides SAP Adaptive Server Enterprise (ASE) users the ability to use SAP HANA on SAP ASE data for real-time analytics.</p> <p>For more information, see the documentation for SAP HANA accelerator for SAP ASE on SAP Help Portal.</p>
SAP HANA remote data sync	hdbrdsyncserver	rdsyncserver	<p>The remote data sync server is part of SAP HANA real-time replication. SAP HANA remote data sync is a session-based synchronization technology designed to synchronize SAP SQL Anywhere remote databases with a consolidated database.</p> <p>For more information, see the documentation for SAP HANA remote data sync on SAP Help Portal.</p>

SAP HANA XS Advanced Services

SAP HANA includes a run-time environment for application development: **SAP HANA extended application services, advanced model** (XS advanced). The SAP HANA XS advanced model represents an evolution of the application server architecture within SAP HANA by building upon the strengths (and expanding the scope) of previous SAP HANA extended application services, classic model (XS classic).

The SAP HANA XS advanced runtime runs either on dedicated hosts or together with other SAP HANA components on the same host. If the runtime platform of SAP HANA XS advanced is installed in your system, the following additional services run in the system database for platform services and for executing applications:

Server Component	OS Process	Service Name	Description
SAP HANA XS Controller	hdbxscontroller	xscontroller	<p>The Controller is the central management component of SAP HANA XS advanced. For example, it has a view on all deployed and/or running applications, and persists configuration and status information in the database.</p> <p>The Platform Router instance is managed by the <code>xscontroller</code> service. The Platform Router, which is realized by an SAP Web Dispatcher instance, exposes the public endpoint for the entire system.</p>

Server Component	OS Process	Service Name	Description
SAP HANA XS Execution Agent	hdbxsexeagent	xsexecagent	The Execution Agent is responsible for managing processes, that is starting, keeping alive, and stopping tasks.
SAP HANA XS User Authentication and Authorization (UAA)	hdixsuaaserver	hdixsuaaserver	The UAA service manages user logon and logoff requests in SAP HANA XS advanced.

→ Recommendation

SAP recommends that customers and partners who want to develop new applications use SAP HANA XS advanced model. If you want to migrate existing SAP HANA XS classic applications to run in the new SAP HANA XS advanced run-time environment, SAP recommends that you first check the features available with the installed version of SAP HANA XS advanced; if the SAP HANA XS advanced features match the requirements of the SAP HANA XS classic application you want to migrate, then you can start the migration process. For more information, see the *SAP HANA XS Advanced Migration Guide*.

Related Information

[Ports and Connections \[page 1043\]](#)

[The JSON Document Store \[page 484\]](#)

[Enabling HDI in the Database \[page 1460\]](#)

[SAP Note 1650957](#)

[SAP Note 2465027](#)

[SAP HANA Dynamic Tiering](#)

[SAP HANA Streaming Analytics](#)

[SAP HANA Accelerator for SAP ASE](#)

[SAP HANA Remote Data Sync](#)

4.5 System Limitations

Limitations to take into consideration when administering an SAP HANA database.

Aside from the table below, most system limits can also be viewed by querying the M_SYSTEM_LIMITS system view (`SELECT * FROM M_SYSTEM_LIMITS;`). However, your values might differ depending on the hardware and software configuration your system uses.

Limitation Area	Limit	M_SYSTEM_LIMITS view name for the limitation
-----------------	-------	--

Database size limit	Row Store: 1,945 GB Column Store: Dependent on size of physical memory	MAXIMUM_SIZE_OF_ROW_STORE
Number of locks	Unlimited for record locks, 16,383 for table locks	MAXIMUM_NUMBER_OF_TABLE_LOCKS
Number of sessions	65,536	MAXIMUM_NUMBER_OF_SESSIONS
Number of tables in row store	1,800,000	

Schema Limitations

Number of schemas per SAP HANA instance	Maximum value of BIGINT data type	
Number of tables in a schema	131,072	MAXIMUM_NUMBER_OF_TABLES_IN_SCHEMA
Identifier length	127 characters	MAXIMUM_LENGTH_OF_IDENTIFIER
Length of an alias name	128 characters	MAXIMUM_LENGTH_OF_ALIAS_NAME
Table name length	Same as Identifier length	MAXIMUM_LENGTH_OF_IDENTIFIER
Column name length	Same as Identifier length	MAXIMUM_LENGTH_OF_IDENTIFIER
Length of a string literal	8 MB	MAXIMUM_LENGTH_OF_STRING_LITERAL
Number of hex characters in a binary literal	8,192 Bytes	MAXIMUM_LENGTH_OF_BINARY_LITERAL

Tables and View Limitations

Number of columns in a table	1,000 This limit can vary based on context, however. For example, in the context of virtual tables, SAP HANA may be limited by the capabilities of the remote system. In the case of virtual tables, for example, the limit of the other DBMS may apply instead. In contexts such as these, the limit that is met first becomes the actual limit.	MAXIMUM_NUMBER_OF_COLUMNS_IN_TABLE
Number of columns in a column store table	64,000	MAXIMUM_NUMBER_OF_COLUMNS_IN_COLUMN_TABLE

Number of columns in a view	64,000	MAXIMUM_NUMBER_OF_COLUMNS_IN_VIEW
Number of rows in each table	Limited by storage size RS: 1,945 GB/sizeof(row), CS: 2 ³¹ * number of partitions	
Length of a row	Limited by RS storage size (1,945 GB per index server)	
Size of a non-partitioned table	Limited by RS storage size (1,945 GB per index server)	
Number of partitions in a CS table	16,000	MAXIMUM_NUMBER_OF_PARTITIONS_IN_CSTABLE
Number of triggers per table per DML statement	1,024	MAXIMUM_NUMBER_OF_TRIGGERS_PER_TABLE_PER_DML
Number of records per (non-partitioned) table	2 ³¹ (i.e. 2,147,384,648, or ~2 billion)	

Indexes and Constraints

Number of indexes for a table	1,023	MAXIMUM_NUMBER_OF_INDEXES_IN_TABLE
Number of primary key columns in each table	16	MAXIMUM_NUMBER_OF_COLUMNS_IN_PRIMARY_KEY
Number of primary key columns in each column store table	1,000	MAXIMUM_NUMBER_OF_COLUMNS_IN_PRIMARY_KEY_IN_COLUMN_TABLE
Number of columns in an index	16	MAXIMUM_NUMBER_OF_COLUMNS_IN_INDEX
Number of columns in a UNIQUE constraint	16	MAXIMUM_NUMBER_OF_COLUMNS_IN_UNIQUE_CONSTRAINT
Size of sum of primary key, index, UNIQUE constraint	16,384 Bytes	MAXIMUM_SIZE_OF_KEY_IN_INDEX
Number of indexes in row store	256,000	

SQL

Length of an SQL statement	2,147,483,648 Bytes	MAXIMUM_LENGTH_OF_SQL_STATEMENT
----------------------------	---------------------	---------------------------------

Depth of SQL view nesting	128	MAXI- MUM_DEPTH_OF_SQL_VIEW_NEST- ING
Depth of SQL parse tree	255	MAXI- MUM_DEPTH_OF_SQL_PARSE_TREE
Number of joined tables in an SQL statement or view	255	MAXIMUM_NUMBER_OF_JOIN_TA- BLES_IN_STATEMENT
Number of columns in an ORDER BY	65,535	MAXIMUM_NUMBER_OF_COL- UMNS_IN_ORDER_BY
Number of columns in a GROUP BY	65,535	MAXIMUM_NUMBER_OF_COL- UMNS_IN_GROUP_BY
Number of elements in IN predicates	65,535	MAXIMUM_NUMBER_OF_COL- UMNS_IN_IN_PREDICATE
Number of elements in SELECT clause	65,535	MAXIMUM_NUMBER_OF_OUT- PUT_COLUMNS_IN_STATEMENT
Number of tables in a statement	4,095	MAXIMUM_NUMBER_OF_TA- BLES_IN_STATEMENT

LOB Limitations

Maximum size of an in-memory LOB for a column store table	1,073,741,821 bytes	MAXIMUM_SIZE_OF_MEM- ORY_LOB_IN_COLUMN_STORE
Maximum size of an in-memory LOB for a row store table	2,147,463,647 bytes	MAXIMUM_SIZE_OF_MEM- ORY_LOB_IN_ROW_STORE
Maximum size of a packed LOB	1,013,760 bytes	MAXIMUM_SIZE_OF_PACKED_LOB
Maximum size of a LOB on disk	4,294,967,295 bytes	MAXIMUM_SIZE_OF_DISK_LOB

Procedures

Size of all stored procedures	1,945 GB	MAXI- MUM_SIZE_OF_ALL_STORED_PROCE- DURES
Size of a procedure definition	2 GB	MAXIMUM_SIZE_OF_PROCE- DURE_DEFINITION

5 SAP HANA Administration Tools

Several tools can be used for the administration of SAP HANA.

Native Tools for SAP HANA Administration

SAP HANA Cockpit

The SAP HANA cockpit provides a single point of access to a range of tools for the administration and detailed monitoring of SAP HANA databases. For example, you can use the cockpit to start and stop systems and databases, monitor databases, configure system settings, manage users and authorizations, and do backups.

The cockpit also integrates the SAP HANA database explorer and the SAP HANA SQL analyzer. The database explorer allows you to query information about the database and view information about catalog objects, while the SQL analyzer helps you to understand and analyze the execution plans of queries.

The SAP HANA cockpit is a Web-based HTML5 user interface that you access through a browser. It runs on SAP HANA extended application services, advanced model (XS advanced). It can be used to administer and monitor SAP HANA databases running SAP HANA 1.0 SPS 12 or higher.

i Note

An SAP HANA cockpit support pack (SP) is released with every SAP HANA platform support package stack (SPS), but additional cockpit SPs may be released between platform SPSs. The *SAP HANA Administration Guide* contains information only about the cockpit SP delivered with the corresponding platform SPS. If you have a more recent cockpit SP, refer to the documentation available at https://help.sap.com/viewer/p/SAP_HANA_COCKPIT. For more information about the revision and maintenance strategy of the cockpit, see SAP Note [2433181](#).

More information

[SAP HANA Cockpit \[page 42\]](#)

[SAP HANA Cockpit Installation and Update Guide](#)

SAP HANA Studio

As an administrator, you use the SAP HANA studio, for example, to start and stop systems, monitor the system, configure system settings, and manage users and authorizations. Developers can use the SAP HANA studio to create content such as modeled views and stored procedures in the SAP HANA repository.

The SAP HANA studio is the legacy development environment and administration tool for SAP HANA, based on the Eclipse platform.

More information

[SAP HANA Studio \[page 113\]](#)

SAP HANA hdbsql

SAP HANA hdbsql allows you to execute SQL statements and database procedures, as well as query information about the database and catalog objects, from the command line.

SAP HANA hdbsql is a command line tool installed on the SAP HANA server. It is available at `/usr/sap/<SID>/HDB<instance>/exe`. It can be used to access databases on both local and remote computers.

More information

[SAP HANA HDBSQL \(Command-Line Reference\) \[page 1963\]](#)

SAP HANA XS Administration Tools

Both the SAP HANA XS classic model and the SAP HANA XS advanced model include Web-based applications to configure and maintain the basic administration-related elements of the application-development process.

The SAP HANA XS advanced cockpit can be accessed from the SAP HANA cockpit.

In addition, the SAP HANA XS advanced model provides command-line tools (`xsa` and `xs CLI`) that you can use to maintain the applications that are deployed to the XS advanced run-time environment, as well as specific elements of the run-time environment itself, for example, the components that enable it, and the users who access and use it.

More information

[Maintaining the SAP HANA XS Classic Model Run Time \[page 1526\]](#)

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1647\]](#)

SAP HANA Application Lifecycle Management

Separate graphical user interfaces are available to install, update, and uninstall products and software components in the SAP HANA XS advanced and the SAP HANA XS classic environments.

The XS Advanced Application Lifecycle Management GUI can be accessed from the SAP HANA cockpit, while the XS Classic Application Lifecycle Management GUI can be accessed from the SAP HANA studio.

In addition, the `hdbalcm` command line tool can be used in the SAP HANA XS classic environment, while the `xs CLI` can be used in the SAP HANA XS advanced environment.

More information

[Installing and Updating SAP HANA Products and Software Components in SAP HANA XS Classic Model \[page 952\]](#)

[Installing and Updating Products and Software Components in SAP HANA XS Advanced Model \[page 960\]](#)

SAP HANA Lifecycle Manager

The SAP HANA database lifecycle manager (HDBLCM) is used to perform SAP HANA platform lifecycle management (LCM) tasks, including installing, updating, and configuring an SAP HANA system.

The SAP HANA HDBLCM program can be run as a graphical user interface, a command-line interface, or as Web user interface in a Web browser.

More information

[About the SAP HANA Database Lifecycle Manager \(HDBLCM\) \[page 921\]](#)

SAP HANA Hardware Configuration Check Tool

The SAP HANA Hardware Configuration Check Tool (HWCCT) is a framework that provides tests and reports for new single node appliances and scale-out systems to determine if the hardware you intend to use meets the minimum performance criteria required to run SAP HANA in production use.

More information

[SAP HANA Hardware Configuration Check Tool for Tailored Data Center Integration \[page 142\]](#)

SAP Tools for SAP HANA Administration

SAP Solution Manager

Use SAP Solution Manager to perform end-to-end root cause analysis and unified alert inbox for entire landscape and business process reporting.

SAP Solution Manager is a central alerting and monitoring infrastructure running on SAP NetWeaver AS for ABAP.

More information

[SAP Solution Manager for SAP HANA Administration \[page 144\]](#)

SAP Landscape Manager

SAP Landscape Manager allows you to automate advanced SAP HANA operations and avoiding business downtime during maintenance activities.

It is a central landscape management solution running SAP NetWeaver AS for Java.

More information

[SAP Landscape Virtualization Management, enterprise edition](#)

SAP IT Operations Management

SAP IT Operations Analytics (ITOA) lets you get and maintain a holistic, real-time overview of complex datacenter landscapes. You can collect, process, and analyze large volumes of data to find the root causes of datacenter issues and resolve them swiftly, or to predict issues and prevent them from happening in the first place.

ITOA is an SAP HANA XSC application that runs on the SAP HANA platform with the SAP HANA Predictive Analysis Library (SAP HANA PAL).

More information

[SAP IT Operations Analytics](#)

5.1 SAP HANA Cockpit

Use the Web-based administration tool SAP HANA cockpit for the administration, monitoring and maintenance of SAP HANA systems.

The SAP HANA cockpit provides tools for the administration and monitoring of SAP HANA databases, and development capabilities through the SAP HANA database explorer. You can manage multiple resources, each running version SAP HANA 1.0 SPS 12, or later. Resources running version SAP HANA 2.0 SPS 01 or later run in multi-container mode, but you can also monitor single-container systems running earlier versions of SAP HANA.

i Note

An SAP HANA cockpit support pack (SP) is released with every SAP HANA platform support package stack (SPS), but additional cockpit SPs may be released between platform SPSs. The *SAP HANA Administration Guide* contains information only about the cockpit SP delivered with the corresponding platform SPS. If you have a more recent cockpit SP, refer to the documentation available at https://help.sap.com/viewer/p/SAP_HANA_COCKPIT. For more information about the revision and maintenance strategy of the cockpit, see SAP Note [2433181](#).

What can I do with the cockpit?

The SAP HANA cockpit provides aggregate, system and database administration features, such as database monitoring, user management, and data backup. You can use the SAP HANA cockpit to start and stop systems or services, monitor the system, configure system settings, and manage users and authorizations.

Cockpit apps that allow you to manage SAP HANA options and capabilities (for example, SAP HANA dynamic tiering) are only available if the option or capability has been installed.

How can I keep an eye on the big picture?

When you first launch the cockpit, you can see system and tenant databases. (The cockpit refers to these as resources). A resource is an SAP HANA system (identified by a host name and instance number) which may be a system or tenant database in a tenant (database) container, or a system in a single database container. These resources are organized into resource groups - you'll only see resources belonging to the groups to which your cockpit user has been granted access. At a glance, you can see top alerts from more than one resource, compare resource configurations and monitor the health of multiple resources.

Whenever you like, you can drill down to perform in-depth monitoring on an individual system or tenant. In order to see alerts and other data for this individual resource you'll need to enter database user credentials. These database user credentials must preexist (i.e. they will have already been created on the resource you are drilling into), and must have the system privilege CATALOG READ and SELECT on _SYS_STATISTICS. For any systems running version SAP HANA 2.0 SPS 01, or later, the cockpit resource administrator has the option to enable or enforce single sign-on (SSO).

How do I get access to groups of resources?

A single `COCKPIT_ADMIN` user is created through the cockpit installation process. This user creates other cockpit users through the Cockpit Manager configuration tool, which is launched through a separate URL provided during installation.

The cockpit administrator assigns the role of `Cockpit Resource Administrator` to at least one cockpit user. The `Cockpit Resource Administrator` registers resources, again through the Cockpit Manager. When resources are registered they are added to auto-generated resource groups, based on system usage type.

Since the `Cockpit Resource Administrator` cannot grant cockpit users access to an auto-generated resource group, they must also create one or more custom resource groups. They add registered resources to each group, and grant access to one or more of the cockpit users which were created by the `COCKPIT_ADMIN`. When you launch the cockpit, you'll be able to see all the registered resources that belong to each of the resource groups to which the `Cockpit Resource Administrator` has granted you access.

Integrated into the cockpit is the SAP HANA database explorer. The database explorer provides the ability to query information about the database using SQL and MDX statements, as well as the ability view information about your database's catalog objects.

Related Information

[SAP Note 2380291](#) 

5.1.1 Set up SAP HANA Cockpit for the First Time

After installation and before other users are able to access the SAP HANA cockpit, you need to perform several steps.

Prerequisites

- You have access to the cockpit administrator (`COCKPIT_ADMIN`) user, created during the installation process
- You have access to the master password which you were prompted to enter during the installation process.
- You know the URLs for the cockpit and the cockpit manager, created during the installation process.

Procedure

1. Connect to the Cockpit Manager and sign in as the `COCKPIT_ADMIN` user. The `COCKPIT_ADMIN` user and corresponding master password were established during cockpit installation.

You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. Create other cockpit users.
3. For each resource you plan to register, create a technical user with CATALOG_READ database privilege and SELECT granted on the _SYS_STATISTICS schema.

i Note

It's not possible by using the cockpit to create the technical user required to register a resource in the SAP HANA cockpit. You need to create this user and grant the minimum necessary authorization by using SQL as follows:

```
CREATE USER <username> PASSWORD <password> NO FORCE_FIRST_PASSWORD_CHANGE;  
GRANT CATALOG READ to <username>;  
GRANT SELECT on SCHEMA _SYS_STATISTICS to <username>
```

4. Register resources.
5. Create resource groups and add registered resources to each group.
6. Assign cockpit users to have access to one or more groups of resources.
7. Determine whether you would like to configure single sign-on authentication to the cockpit and/or the registered resources.
8. Share the credentials of the newly-created cockpit users with the appropriate people, and instruct them to sign in to the SAP HANA cockpit .

In the cockpit, the [Resource Directory](#) displays all the registered resources to which this cockpit user has access.

Related Information

[Open SAP HANA Cockpit \[page 47\]](#)

[Working with Resources and Resource Groups \[page 69\]](#)

[Setting Up Single Sign-On \[page 88\]](#)

[Managing Cockpit Users \[page 60\]](#)

[SAP Note 2380291](#) 

5.1.1.1 Determine Ports for SAP HANA Cockpit and Cockpit Manager

The ports for SAP HANA cockpit and the cockpit manager can be determined after the installation.

Prerequisites

- You are logged in as `<sid>adm` user.
- You know the XS organization manager user password. The password matches the master password, which is set during installation.

Context

Ports, through which the SAP HANA cockpit and the cockpit manager can be accessed, are assigned automatically by the installer. Once the cockpit installation is successfully completed, information about host and ports is displayed. If this information is no longer available, you can execute the following commands in the XS console to determine ports.

You can also assign free ports to SAP HANA cockpit during installation. For more information, see *SAP Note 2389709* in Related Information.

Procedure

1. Change to the directory that contains the XS Advanced installation:

```
cd <sapmnt>/<SID>/xs/bin
```

By default, `<sapmnt>` is `/hana/shared`.

2. Log on to the SAP HANA XS advanced runtime. To do this, use the following command:

```
./xs-admin-login
```

3. Enter the XS organization manager user password.
4. Display a list of the applications running in the current space. In the command shell, run the following command:

```
xs apps
```

A list of all running apps is displayed. Information on host and ports are displayed in the `urls` column. The SAP HANA cockpit is listed as `cockpit-web-app`. The cockpit manager is listed as `cockpit-admin-web-app`.

Output Code

```
Getting apps in org "HANACockpit" / space "SAP" as COCKPIT_ADMIN...
Found apps:
name                requested state  instances  memory
disk                urls
-----
auditlog-db         STOPPED         0/1        16.0 MB
<unlimited>         <none>
auditlog-server     STARTED         1/1        256 MB
<unlimited>         https://<hostname>:51002
auditlog-broker     STARTED         1/1        64.0 MB
<unlimited>         https://<hostname>:51003
deploy-service      STARTED         1/1        280 MB
<unlimited>         https://<hostname>:51004
auditlog-odata      STARTED         1/1        128 MB
<unlimited>         https://<hostname>:51005
component-registry-db STOPPED         0/1        16.0 MB
<unlimited>         <none>
auditlog-ui         STARTED         1/1        64.0 MB
<unlimited>         https://<hostname>:51007
product-installer   STARTED         1/1        256 MB
<unlimited>         https://<hostname>:51006
hrtt-service        STARTED         1/1        512 MB
<unlimited>         https://<hostname>:51009
sqlanz-svc          STARTED         1/1        256 MB
<unlimited>         https://<hostname>:51010
sqlanz-ui           STARTED         1/1        128 MB
<unlimited>         https://<hostname>:51011
hrtt-core           STARTED         1/1        512 MB
<unlimited>         https://<hostname>:51012
sapui5_fesv2        STARTED         1/1        256 MB
<unlimited>         https://<hostname>:51015
sapui5_fesv3        STARTED         1/1        256 MB
<unlimited>         https://<hostname>:51025
cockpit-adminui-svc STARTED         1/1        128 MB
<unlimited>         https://<hostname>:51022
cockpit-collection-svc STARTED         1/1        768 MB
<unlimited>         https://<hostname>:51016
cockpit-hdb-svc     STARTED         1/1        768 MB
<unlimited>         https://<hostname>:51018
cockpit-hdbui-svc   STARTED         1/1        128 MB
<unlimited>         https://<hostname>:51020
cockpit-landscape-svc STARTED         1/1        128 MB
<unlimited>         https://<hostname>:51019
cockpit-persistence-svc STARTED         1/1        768 MB
<unlimited>         https://<hostname>:51017
cockpit-telemetry-svc STARTED         1/1        768 MB
<unlimited>         https://<hostname>:51026
cockpit-xsa-svc     STARTED         1/1        768 MB
<unlimited>         https://<hostname>:51024
cockpit-admin-web-app STARTED 1/1 128 MB
<unlimited>         https://<hostname>:51023
cockpit-web-app STARTED 1/1 512 MB
<unlimited>         https://<hostname>:51021
```

Related Information

[SAP Note 2389709](#)

5.1.2 Open SAP HANA Cockpit

You access the SAP HANA cockpit from a Web browser.

Prerequisites

- You know the URL for the cockpit, created during the installation process.
- You have a cockpit user name and password.
- For each resource you have been authorized to monitor through the cockpit, you have a database user name and password.
- Your Web browser supports the SAPUI5 library `sap.m`.
For more information about SAPUI5 browser support, see SAP Note 1716423 and the Product Availability Matrix (PAM) for SAPUI5.

Procedure

1. Enter the SAP HANA cockpit URL in your browser.
2. Enter your cockpit user name and password.
The SAP HANA cockpit opens.
3. Here you can see all the resources to which you are authorized.
4. For each resource, set your database credentials.
5. Double-click on a specific resource to access the [Overview](#) for that resource.

Related Information

[Monitoring in SAP HANA Cockpit \[page 317\]](#)

[Security Aspects of SAP HANA Cockpit \[page 86\]](#)

[SAP Note 1716423](#) 

[Product Availability Matrix \(PAM\) for SAPUI5](#) 

5.1.3 Authorizations Needed for Monitoring and Administration

To view information about the SAP HANA database and access the various applications for administration and monitoring, you need to connect to the resource with a database user with appropriate database privileges.

i Note

To be able to connect to a resource and see minimum monitoring information, the connecting database user must have system privilege CATALOG READ and the SELECT privilege on the schema `_SYS_STATISTICS`.

Database Monitoring and Administration

These tables list the database privileges required to view information about an SAP HANA database on the [System Overview](#) page and to access monitoring and administration functions on subsequent pages.

Monitoring

To Access...	You Need These SAP HANA Privileges...
Overall Database Status	SELECT privileges on: <ul style="list-style-type: none"> • SYS_DATABASES.m_services • SYS_DATABASES.m_service_memory • SYS_DATABASES.m_service_statistics • SYS_DATABASES.m_heap_memory_reset • SYS.m_services, SYS.m_services • SYS.m_service_memory • SYS.m_service_statistics • SYS.m_heap_memory_reset • _SYS_STATISTICS.STATISTICS_SCHEDULE • _SYS_STATISTICS.HELPER_ALERT_CHECK_INACTIVE_SERVICES_AGE
Alerts	SELECT privileges on <code>_SYS_STATISTICS</code>
Memory Usage	PUBLIC role (no additional authorization required)
Memory Analysis	<ul style="list-style-type: none"> • CATALOG_READ system privilege • SELECT on <code>_SYS_STATISTICS</code> object privileges
CPU Usage	No additional authorization required
Disk Usage	No additional authorization required
Performance Monitor	No additional authorization required
Monitor Statements	No additional authorization required
Sessions	No additional authorization required
Threads	No additional authorization required

To Access...	You Need These SAP HANA Privileges...
Monitor expensive statements	No additional authorization required
Open SQL plan cache	No additional authorization required
Open blocked transactions	No additional authorization required
Database Administration	
To Access...	You Need These SAP HANA Privileges...
Copy database	BACKUP ADMIN
Manage database backups	BACKUP ADMIN
Recover database	Tenant database: DATABASE ADMIN System database: Operating system user <sid>adm
Manage system licenses	System privilege LICENSE ADMIN
Manage statement hints	No additional authorization required
Alerting and Diagnostics	
To Access...	You Need These SAP HANA Privileges...
Configure alerts	SELECT privileges on _SYS_STATISTICS
Plan trace	No additional authorization required
System Replication	
To Access...	You Need These SAP HANA Privileges...
System Replication	No additional authorization required
Other Administration	
To Access...	You Need These SAP HANA Privileges...
Administer XS advanced	Not applicable.
Launch cockpit for XS advanced	For more information about the availability of these tools, see the section <i>Maintaining the XS Advanced Runtime Environment with SAP HANA Advanced Cockpit</i> in the <i>SAP HANA Administration Guide</i> .
Application and Platform Lifecycle Management	
To Access...	You Need These SAP HANA Privileges...
Platform Lifecycle Management	<sid>adm privileges
Application Lifecycle Management	Not applicable. For more information about the availability of the application lifecycle management GUI, see the section on SAP HANA application lifecycle management in the <i>SAP HANA Administration Guide</i> .

General Information

To Access...	You Need These SAP HANA Privileges...
General Information	System privilege CATALOG READ and SELECT on _SYS_STATISTICS
Help links	No additional authorization required

Security

Security

To Access...	You Need These SAP HANA Privileges...
Data Encryption	<ul style="list-style-type: none">• System privilege CATALOG READ• To enable/disable encryption: ENCRYPTION ROOT KEY ADMIN• To view SSFS master key information: RESOURCE ADMIN
Auditing	<p>To see information about auditing status audit policies, you need the system privileges AUDIT ADMIN or CATALOG READ.</p> <p>To see information about audit trail targets, you need the system privilege INIFILE ADMIN.</p>
Authentication	<p>System privilege CATALOG READ</p> <p>To be able to see the password blacklist on opening the Password Policy and Blacklist page, you need SELECT privilege on _SYS_SYS_PASSWORD_BLACKLIST (_SYS_SECURITY).</p>

User & Role Management

To Access...	You Need These SAP HANA Privileges...
Manage users	System privilege CATALOG READ to view users and USER ADMIN to create and manage users
Assign roles to users	System privilege CATALOG READ to view roles and system privilege ROLE ADMIN to assign roles
Assign privileges to users	System privilege CATALOG READ to view users; to assign privileges system privilege USER ADMIN and the privileges required to grant specific privileges to the user
Manage roles	System privilege CATALOG READ to view roles; to create and manage roles system privilege ROLE ADMIN and the privileges required to grant specific privileges to roles

Security Related Links

To Access...	You Need These SAP HANA Privileges...
Manage certificates	One of these system privileges: <ul style="list-style-type: none"> • CATALOG READ • SSL ADMIN • USER ADMIN • TRUST ADMIN • CERTIFICATE ADMIN
Manage certificate collections	One of these system privileges: <ul style="list-style-type: none"> • CATALOG READ • SSL ADMIN • USER ADMIN • TRUST ADMIN • CERTIFICATE ADMIN
Network security information	System privilege CATALOG READ
View anonymization report	System privilege CATALOG READ
Manage SAML identity providers	System privilege USER ADMIN
Security administration help	No additional authorization required
SAP HANA security website	No additional authorization required
Security checklists	No additional authorization required

Performance Management

Workload Management

To Access...	You Need These SAP HANA Privileges...
Analyze Workload (Based on Thread Samples)	System privileges CATALOG READ and INIFILE ADMIN.
Analyze Workload (Based on Engine Instrumentation)	System privilege WORKLOAD ANALYZE ADMIN.
Capture Workload	<p>To capture workloads, you need the WORKLOAD CAPTURE ADMIN privilege.</p> <p>To see the previously used optional filters on the capture configuration page, you need the INIFILE ADMIN privilege.</p> <p>To make a backup from the database, you need the BACKUP OPERATOR privilege.</p>
Replay Workload	<p>To preprocess and replay workloads, you need the WORKLOAD REPLAY ADMIN privilege.</p> <p>To view the load chart in the replay report, you need the WORKLOAD REPLAY ADMIN and the WORKLOAD ANALYZE ADMIN privileges.</p>

5.1.4 Open the SAP HANA Database Explorer (SAP HANA Cockpit)

Open the SAP HANA database explorer to view information about your database's catalog objects and execute queries.

Prerequisites

You must be a user of the SAP HANA cockpit.

Context

The database explorer displays your database objects, grouped by schema. It also contains an SQL console for executing SQL statements and an SQL analyzer for analyzing them, an SQL debugger for debugging procedures, as well as an MDX console for executing MDX queries.

Procedure

1. In the SAP HANA cockpit, navigate to the [Overview](#) page of a database that you want to explore, and then click the [Browse Database](#) link on that page.

The database explorer opens and the catalog browser lists your database in its tree.

You can add other databases by clicking [Add a database to the Database Explorer \(+\)](#) at the top of the catalog browser.

2. From the catalog browser, choose your database to view the object types available, and then choose an object type to view its individual objects.
3. Choose an object to view the object's definition in the right pane, or right-click the object to choose a different action.

Related Information

[Add SAP HANA Cockpit Resources and Databases to the SAP HANA Database Explorer \[page 53\]](#)

[Execute SQL Statements \[page 493\]](#)

5.1.4.1 Add SAP HANA Cockpit Resources and Databases to the SAP HANA Database Explorer

Add a cockpit resource or database to the SAP HANA database explorer so that you can browse its catalog and execute SQL statements against it.

Prerequisites

To add an SAP HANA database to the database explorer, you must have a user name and password for the database.

You must have a user ID and password for the cockpit.

Context

Adding a database to the database explorer is similar to registering a resource in the cockpit. Once a database is added, it is listed in the catalog browser pane on the left.

You cannot add a database to the database explorer that uses LDAP authentication.

Procedure

1. Open the database explorer from the SAP HANA cockpit.
2. Add a database by clicking the [Add a database to the Database Explorer](#) icon (+) at the top of the catalog browser pane on the left.
3. From the [Database Type](#) dropdown list, choose the type of database to add:

Database Type	Action
An SAP HANA cockpit resource database	<ol style="list-style-type: none">1. Choose Cockpit Resource.2. Choose a resource from the list.
An SAP HANA database	<ol style="list-style-type: none">1. Choose SAP HANA database.2. Specify the fully qualified domain name (FQDN) of the host on which the system is installed. Specify the instance number of the database you are adding. When adding a database that is part of a multi-host system, specify the master host. You do not have to enter all host names explicitly as they are determined automatically. If the master host becomes unavailable, then the connection is automatically established through one of the other hosts. Hosts that are added to the system later are also detected automatically.

Database Type	Action
A tenant or the system database that is part of a multitenant database container system	<ol style="list-style-type: none"> 1. Choose <i>SAP HANA database (Multitenant)</i>. 2. Specify the host and instance number of the system database. 3. Specify whether you are adding the system database or a tenant database. If you are adding a tenant database, then specify the name of the tenant database.

4. (Optional) If you are adding an SAP HANA database, a tenant database, or a system database that is part of a multi-tenant database container system, then specify encryption information and advanced connection properties as required.

Choose from the following encryption options:

Option	Description
<i>Save user and password (stored in the SAP HANA secure store.)</i>	<p>By default, your database credentials are not saved. Each time you need to connect to an added database, you must provide your user credentials.</p> <p>Choose this option to save these user credentials so that you do not have to re-enter them each time you connect. These credentials are saved to the SAP HANA secure store.</p>
<i>Connect to the database securely using TLS/SSL. (Prevents data eavesdropping.)</i>	<p>Choose this option to encrypt communication between the database explorer and the SAP HANA database using the Transport Security Layer (TLS)/Secure Sockets Layer (SSL) protocol.</p>
<i>Verify the server's certificate using the trusted certificate below.</i>	<p>Choose the <i>Verify the server's certificate using the trusted certificate below</i> option and provide a trusted certificate, if you want to verify the server's certificate when connecting. This prevents server impersonation. The certificate field must contain the contents of a certificate, and not a file name. For more information, see the <i>SAP HANA Security Guide</i>.</p>

i Note

You can only choose this option if you have also chosen the *Connect to the database securely using TLS/SSL. (Prevents data eavesdropping.)* option.

Choose from the following advanced options.

In the *Advanced Options* field, specify the advanced options as comma-separated name=value pairs. For example: `locale=en-US, isolationLevel=READ COMMITTED`.

Option	Description
autoCommit	Specify ON to have the database explorer implicitly commit your transaction after executing each DDL statement. Otherwise, specify OFF to have your transactions committed after executing a commit statement, and to support rollbacks of DDL statements. The default value is OFF.
isolationLevel	The isolation level for the connection. The supported values are: READ COMMITTED, REPEATABLE READ, and SERIALIZABLE. The default is to specify no isolation level.
locale	The locale to use for the connection. If you do not set this option, then the database explorer looks for a locale setting in your user parameter, and then in your browser. If no locale setting is found, then the locale is set to en-US.
schema	The name of the schema that you want to use
CLIENT	Sets the session client for the connection. The value is a three character string. For example, CLIENT=100.

- (Optional) Choose a display name that is used to identify your database in the database explorer. Each database must have a unique display name. If you do not choose a display name, then one is generated for you.
- Click *OK*.

Results

The resource or database is added to the database browser.

Related Information

[Open Catalog Objects \[page 56\]](#)

[Execute SQL Statements \[page 493\]](#)

5.1.4.2 Open Catalog Objects

Browse your database's catalog using the SAP HANA database explorer.

Prerequisites

You must be a user of the database that you want to explore and you must have the required privileges to view the catalog items.

Context

Some monitoring and problem analysis may require you to examine individual tables and views, for example, system views provided by the SAP HANA database. Use the catalog browser, which is located in the left pane, to find and open these catalog objects.

Procedure

1. In the catalog browser, choose the database that you want to explore.
If your database is not listed in the catalog browser, then click [Add a database to the Database Explorer \(+\)](#) to add the database.
The catalog browser lists the catalog objects, grouped by schema.
2. Choose an object type to view its objects.
For example, choose [Tables](#) to list the tables in the database.
3. Choose an object to view its definition in an editor in the right pane, or right-click the object to choose a different action.
For example, right-click a table and choose [Open Data](#) to view the table's data.

Related Information

[Add SAP HANA Cockpit Resources and Databases to the SAP HANA Database Explorer \[page 53\]](#)

5.1.4.3 Search for Database Objects in the SAP HANA Database Explorer

Search for database objects across all databases that the SAP HANA database explorer is connected to.

Prerequisites

In the database explorer, connect to each database that you want to search.

Context

Use the catalog browser to find your object when you know the type of object and the database that it exists in.

If you are unsure of the object type or the database that it exists in, then use the database object search. This search looks for matches in either a specified database or all of the connected databases in the database explorer and it can search across more than one object type. For example, use this search to find all tables and procedures that contain the word **production** in their names.

Procedure

From the right sidebar of the database explorer, click *Object Search* (Q).

- a. Specify a search term that is longer than one character.
Specify * to return all results.
- b. Optionally, restrict your search to a specific database.
- c. Target your search to specific object types by selecting them at the bottom of the pane.

Results

The search results appear in the *Search for Database Objects* pane.

Related Information

[Create URL Shortcuts to Database Objects in the SAP HANA Database Explorer \[page 58\]](#)

[Open Catalog Objects \[page 56\]](#)

5.1.4.4 Create URL Shortcuts to Database Objects in the SAP HANA Database Explorer

Use the SAP HANA database explorer to create a URL shortcut, and then use it to create a bookmark or favorite in your browser.

Procedure

1. In the catalog browser, right-click a database object and choose *Create Shortcut*.
2. Click *Copy URL* to copy the shortcut to your clipboard.
3. Paste the URL into your browser.
4. Use your browser to create a bookmark or favorite for the copied URL.

Results

The shortcut is created. Use it to open the database explorer and quickly navigate to the database object.

Related Information

[Search for Database Objects in the SAP HANA Database Explorer \[page 57\]](#)

5.1.4.5 About the SAP HANA Database Explorer and the SQL Analyzer

Use the SAP HANA database explorer to query information about the database, as well as view information about your database's catalog objects. Use the SAP HANA SQL analyzer to understand and analyze the execution plans of your queries.

The database explorer is integrated into both the SAP Web IDE for SAP HANA and in the SAP HANA cockpit. The database explorer contains features and functions required by both database administrators and developers. For example:

- | | |
|--------------------------|---|
| A catalog browser | View the definitions of all types of catalog objects, for example: tables, views, stored procedures, functions, and synonyms. Also, view the content (data) of your tables and views. |
| An SQL console | Create SQLScript procedures and queries, and then execute them or analyze their performance using the SQL analyzer. |

An SQL analyzer	View detailed information on your queries and evaluate potential bottlenecks and optimizations for these queries. The SQL analyzer is accessible from the SQL console, as well as from the plan trace and expensive statement features in the SAP HANA cockpit.
An MDX console	Create and run MDX queries.
An SQL debugger (SAP Web IDE for SAP HANA)	View the call stack, set break points, view and evaluate expressions and variables. This feature is available only for procedures in HDI containers.

Related Information

[Open the SAP HANA Database Explorer \(SAP HANA Cockpit\) \[page 52\]](#)

[Analyzing Statement Performance \[page 440\]](#)

[SAP Note 2373065](#)

5.1.5 Setup and Administration with the Cockpit Manager

The Cockpit Manager configuration tool is an application separate from the SAP HANA cockpit itself. Launch the Cockpit Manager using the URL that was provided during cockpit installation.

The functionality visible in the Cockpit Manager depends on the role(s) assigned to the user accessing the Cockpit Manager. The administrator roles of cockpit administrator, cockpit user administrator and cockpit resource administrator can perform the tasks necessary to enable users of SAP HANA cockpit to manage and monitor resources. These roles can be assigned together to one user, or to separate individuals.

Before other cockpit users can make full use of SAP HANA cockpit, the cockpit user administrator needs to use the Cockpit Manager to:

- Create and manage cockpit users

Then, the cockpit resource administrator needs to use the Cockpit Manager to:

- Register resources
- Create resource groups
- Add resources to resource groups
- Grant cockpit users access to resource groups

Optionally, the cockpit administrator can to use the Cockpit Manager to:

- Modify settings related to the configuration of the cockpit.

Note

During cockpit installation, a master user, COCKPIT_ADMIN, is automatically created. Its password corresponds to the master password which you were prompted to enter during the installation process. This master user is assigned all three administrator roles, and can therefore access all aspects of the Cockpit Manager, and can create users, register resources, and assign users and resources to resource groups. However, you may wish to assign administrator roles to other users.

Related Information

- [Set up SAP HANA Cockpit for the First Time \[page 43\]](#)
- [Determine Ports for SAP HANA Cockpit and Cockpit Manager \[page 45\]](#)
- [Open SAP HANA Cockpit \[page 47\]](#)
- [Managing Cockpit Users \[page 60\]](#)
- [Working with Resources and Resource Groups \[page 69\]](#)
- [Configuring Cockpit Settings \[page 67\]](#)
- [Security Aspects of SAP HANA Cockpit \[page 86\]](#)
- [Using XS CLI Commands to Troubleshoot the Cockpit \[page 112\]](#)

5.1.5.1 Managing Cockpit Users

Using the Cockpit Manager administration tool, the cockpit user administrator, can create cockpit user log on credentials.

A cockpit user administrator is the only role that can create other cockpit users (application users).

Cockpit Users vs. Database Users

Selecting *Cockpit Users* in the Cockpit Manager allows you to create cockpit users that can be assigned to a group of resources. Once cockpit users have access to a resource group, they will be able to monitor each of the resources within the group, as well as see aggregate data for the group. A cockpit user who isn't given access to any resource groups can log in to the cockpit, but their *My Resources* page displays zero resources. A cockpit user who is given access only to a group that contains no resources also sees a display of zero resources.

i Note

Cockpit users are distinct from the database user credentials associated with the individual resources managed and monitored through the cockpit. Before they can drill down from the aggregate monitoring in the cockpit to a specific system, each cockpit user is required to enter existing database user credentials (unless single sign-on is in effect). Database users are not managed through the Cockpit Manager, but rather through the *Manage Users* link on the system *Overview*.

Each cockpit user must be assigned at least one cockpit role, which will dictate what portions of the cockpit or the Cockpit Manager they can access. (Cockpit roles are unrelated to the roles associated with database users. The latter govern which SAP HANA privileges are assigned to a database user).

Cockpit Role	Permits access to
Cockpit Administrator	The <i>Cockpit Settings</i> section of the Cockpit Manager, where they can configure cockpit settings.

Cockpit Role	Permits access to
Cockpit Resource Administrator	The <i>Registered Resource</i> and <i>Resource Groups</i> sections of the Cockpit Manager, where they can register resources, create resource groups, and assign cockpit users and registered resources to resource groups.
Cockpit User Administrator	The <i>Manage Users</i> section of the Cockpit Manager, where they can create and manage cockpit users.
Cockpit User	The SAP HANA cockpit, where they can view the resources in the resource groups to which they have been granted access.
Cockpit Power User	The SAP HANA cockpit, and the <i>Registered Resource</i> section of the Cockpit Manager.

A user who has only the `Cockpit User` role cannot access the Cockpit Manager. Conversely, a user who has only the `Cockpit User Administrator` or the `Cockpit Resource Administrator` role cannot access the cockpit.

Note

During cockpit installation, a master user, `COCKPIT_ADMIN`, is automatically created. Its password corresponds to the master password which you were prompted to enter during the installation process. This master user is assigned all three administrator roles, and can therefore access all aspects of the Cockpit Manager, and can create users, register resources, and assign users and resources to resource groups. However, you may wish to assign administrator roles to other users.

Because cockpit roles were introduced as of SAP HANA 2.0 SPS 01, if your `COCKPIT_ADMIN` was created during the installation of an earlier version, you may wish to assign it additional roles if you want to continue to access all aspects of the cockpit and the Cockpit Manager, and you will have to log on again to have this change take effect.

Related Information

[Create or Enable a Cockpit User \[page 62\]](#)

[Grant a Cockpit User Access to Specific Groups \[page 63\]](#)

[Add or Remove Users in Resource Groups \[page 84\]](#)

[Managing Resource Groups \[page 80\]](#)

[Register a Resource \[page 70\]](#)

[Security Aspects of SAP HANA Cockpit \[page 86\]](#)

5.1.5.1.1 Create or Enable a Cockpit User

As a cockpit user administrator, you can create new cockpit users, or allow existing business users to access the SAP HANA cockpit.

Context

The *Create User* wizard allows you to create cockpit users by entering new credentials or by choosing to provide cockpit access to existing business users that have been created through other means.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user administrator.
You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the *Manage Cockpit* link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```
2. In the Cockpit Manager, select *Cockpit Users*.
3. Select *Create User*.
4. In the Create User wizard, do one of the following:
 - Select the checkbox to enable an existing business user to access the cockpit, then select the user from the drop-down list.
 - Enter a new user name, enter and confirm a password, and enter an e-mail address.
5. Assign a cockpit role to the user. Choose one, or a combination of:

Cockpit Role	Permits access to
Cockpit Administrator	The <i>Cockpit Settings</i> section of the Cockpit Manager, where they can configure cockpit settings.
Cockpit Resource Administrator	The <i>Registered Resource</i> and <i>Resource Groups</i> sections of the Cockpit Manager, where they can register resources, create resource groups, and assign cockpit users to resource groups.
Cockpit User Administrator	The <i>Manage Users</i> section of the Cockpit Manager, where they can create and manage cockpit users.
Cockpit User	The SAP HANA cockpit, where they can view all resources in any assigned resource groups.
Cockpit Power User	The SAP HANA cockpit, and the <i>Registered Resource</i> section of the Cockpit Manager.

i Note

The Cockpit Power User role is assigned by selecting the checkbox *Allow this user to register resources*.

6. If the cockpit user administrator that you used to access the Cockpit Manager has also been assigned the `Cockpit Resource Administrator` role, you can choose [Add Groups](#) and select the group(s) that the users which you create will be able to access. (Any `Cockpit Resource Administrator` can also do this through the [Resource Groups](#) section of the Cockpit Manager.)
7. (Optional) Scroll up and review what you entered.
8. Select [Create User](#).

On the Cockpit Users page, you can see the list of cockpit users, which includes the one you just added or enabled.

i Note

A new Cockpit User will be able to access the cockpit, but they will not see any resources until someone with the `Cockpit Resource Administrator` role assigns them to at least one resource group that contains at least one registered resource.

Related Information

[Grant a Cockpit User Access to Specific Groups \[page 63\]](#)

[Edit Settings for a Cockpit User \[page 64\]](#)

5.1.5.1.2 Grant a Cockpit User Access to Specific Groups

Add or remove resource groups from the list of groups that a specific user can access.

Context

In order for cockpit users to monitor and manage registered resources, you need to assign them to the resource group to which these resources belong.

i Note

This task can be done in the [Resource Groups](#) section of the Cockpit Manager by a Cockpit Resource Administrator. However, if you, as the Cockpit User Administrator also have the `Cockpit Resource Administrator` role, you have the option to make these changes through the [Cockpit Users](#) section.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user administrator with cockpit resource administrator privileges (a user with the roles `COCKPIT_USER_ADMIN` and `COCKPIT_RESOURCE_ADMIN`).

You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. In the Cockpit Manager, select [Cockpit Users](#).
3. In the left pane, select the user whose access you want to modify.
4. Select [Grant Access to Resource Groups](#). (This link is only visible if the Cockpit User Administrator has been assigned the Cockpit Resource Administrator role.)
5. Select the group(s) that the user will be permitted to access and click [OK](#).

The group(s) to which the user now has access are listed under [User Details](#).

6. (Optional) To remove the user's access to a group, click the red ⊗ to the right of the group that you want to remove from this user, and confirm the request.

Results

The group is no longer listed under [User Details](#). The group can be reassigned to this user or other users at any time.

Related Information

[Add or Remove Users in Resource Groups \[page 84\]](#)

[Create or Enable a Cockpit User \[page 62\]](#)

[Delete a Cockpit User or Revoke Cockpit Access \[page 65\]](#)

5.1.5.1.3 Edit Settings for a Cockpit User

As the cockpit user administrator, you can modify some settings for cockpit users.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user administrator.
You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. In the Cockpit Manager, select [Cockpit Users](#).
3. In the left pane, select the user whose settings you want to modify.

4. Click *Edit*.
5. Assign a cockpit role to the user. Choose one, or a combination of:

Cockpit Role	Permits access to
Cockpit Administrator	The <i>Cockpit Settings</i> section of the Cockpit Manager, where they can configure cockpit settings.
Cockpit Resource Administrator	The <i>Registered Resource</i> and <i>Resource Groups</i> sections of the Cockpit Manager, where they can register resources, create resource groups, and assign cockpit users to resource groups.
Cockpit User Administrator	The <i>Manage Users</i> section of the Cockpit Manager, where they can create and manage cockpit users.
Cockpit User	The SAP HANA cockpit, where they can view all resources in any assigned resource groups.
Cockpit Power User	The SAP HANA cockpit, and the <i>Registered Resource</i> section of the Cockpit Manager.

i Note

The Cockpit Power User role is assigned by selecting the checkbox *Allow this user to register resources*.

6. Click *Apply*.

Related Information

[Grant a Cockpit User Access to Specific Groups \[page 63\]](#)

[Delete a Cockpit User or Revoke Cockpit Access \[page 65\]](#)

5.1.5.1.4 Delete a Cockpit User or Revoke Cockpit Access

As the cockpit user administrator, you can delete cockpit users. However, you will need to decide whether the underlying business user should remain.

Context

Some cockpit users may have been created outside of the cockpit manager; their original purpose may have been to access other applications. Other cockpit users have been created through the cockpit manager for the sole purpose of accessing the cockpit. When you delete a cockpit user, you can choose to:

- Allow the underlying business user to remain (so that it can be used to access other applications), and only revoke the access to the cockpit.
- Completely delete the user.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user administrator.
You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. In the Cockpit Manager, select [Cockpit Users](#).
3. In the left pane, select the user.
4. Select [Delete](#).
5. In the [Confirm Request](#) dialog:

Choose...	In order to...
Remove Access Only	Allow the business user to continue to exist outside the cockpit, but without cockpit access
Delete User	Delete the business user from the cockpit and from any additional applications
Cancel	Neither delete the user nor the cockpit access.

Results

When you choose to remove access or completely delete the user, the user is removed from the list on the Cockpit Users page.

Related Information

[Create or Enable a Cockpit User \[page 62\]](#)

[Grant a Cockpit User Access to Specific Groups \[page 63\]](#)

[Security Aspects of SAP HANA Cockpit \[page 86\]](#)

5.1.5.2 Configuring Cockpit Settings

In the Cockpit Manager, as a cockpit administrator, you can select [Settings](#) to configure data collection, proxy server settings for e-mailed alerts, and the connection timeout period, and to control whether or not SAP HANA Cockpit displays auto-created groups.

Related Information

[Setting Data Collection \[page 67\]](#)

[Setting Proxy Server \[page 68\]](#)

[Setting Connection Timeout \[page 68\]](#)

[Specifying Display of Auto-Generated Groups \[page 69\]](#)

[Working with Resources and Resource Groups \[page 69\]](#)

5.1.5.2.1 Setting Data Collection

As a cockpit administrator, you can reconfigure the default, preconfigured SAP HANA cockpit global settings for collecting monitoring data, such as system status, alert counts, and other data from registered resources.

In the *Cockpit Manager*, select [Cockpit Settings](#), and then [Data Collection](#).

If necessary, you can change the defaults to specify:

- How many worker threads the collection service should use. Increasing threads can improve response time but uses more memory. The default is 5 threads.
- Whether and how often the cockpit collects system status and alert counts. The default is 60 seconds.
- Whether and how often the cockpit collects key performance area monitoring data from each managed resource. The default is 5 minutes.

There may be a brief lag before your changes in values take effect.

i Note

You can also modify the collection settings for a specific resource by editing the details of that resource. Doing so overrides the global settings for that particular resource. See [Override Data Collection for a Resource](#).

→ Tip

The cockpit can support 1000s of registered resources. If the System Health Monitor displays 'Not Collected' for specific resources, you may wish to investigate the collection service log for rejected collections and reconfigure the worker threads accordingly.

Related Information

[Using XS CLI Commands to Troubleshoot the Cockpit \[page 112\]](#)

5.1.5.2.2 Setting Proxy Server

As a cockpit administrator, you can optionally set up a proxy server to use with SAP HANA cockpit through the *Cockpit Manager* by selecting *Settings*, then *Proxy*.

There are two types of proxies available: the Network proxy and the HTTP(S) proxy. In both cases, you need to specify the host and port number.

For an HTTP(S) proxy you can also specify exceptions that should not use the proxy host. Use the *No Proxy Host* field to enter the exceptions (addresses beginning with the strings you enter, separated by semi-colons).

→ Tip

After setting up a proxy server, be sure to check *Enable* before selecting *Save*.

Related Information

[Setting Data Collection \[page 67\]](#)

[Setting Connection Timeout \[page 68\]](#)

5.1.5.2.3 Setting Connection Timeout

As a cockpit administrator, you can specify the length of time that the SAP HANA cockpit waits for a connection before initiating a timeout through the *Cockpit Manager* by selecting *Settings*, then *Connections*.

If a server connection is unresponsive, you may want to ensure that the cockpit doesn't wait for a response indefinitely. You can configure a standard timeout period, and a timeout period for long running tasks.

Related Information

[Setting Data Collection \[page 67\]](#)

[Setting Proxy Server \[page 68\]](#)

5.1.5.2.4 Specifying Display of Auto-Generated Groups

As a cockpit administrator, you can choose whether or not SAP HANA cockpit displays resources as part of auto-created resource groups, or solely as part of resource groups that you create.

A resource group—a named set of one or more registered resources—controls management and monitoring privileges. When you assign a cockpit user to a resource group that you have created, you enable the user to monitor and manage the group's resources through the cockpit. Each registered resource also belongs to a usage type resource group. These auto-created groups of resources (Production, Test, Development) are based on the system usage type of each resource. System usage type is configured during system installation, or later using the `global.ini` file with the `usage` parameter in the `system_information` section.

You can choose to hide one or more of the auto-created groups through the *Cockpit Manager* by selecting *Settings*, then *Display*, and deselecting or selecting each box. Opting to hide the auto-created groups does not affect the system usage type associated with the resource. It simply prevents the cockpit from organizing the display of resources by auto-created group.

Related Information

[Working with Resources and Resource Groups \[page 69\]](#)

5.1.5.3 Working with Resources and Resource Groups

Through the *Cockpit Manager*, any user with the `Cockpit Resource Administrator` role can register resources and create groups of resources that other cockpit users will be able to access with SAP HANA cockpit.

SAP HANA cockpit provides aggregate monitoring, which means you and other cockpit users can see data from multiple resources simultaneously. Resource groups govern which resources each user is permitted to monitor and manage.

Selecting *Registered Resources* allows you to see all the resources that have been registered, and to register more.

Selecting *Resource Groups* allows you to see all the groups that have been created, along with the resources and cockpit users that have been associated with each. Here, you can also create a new resource group, or add a resource or cockpit user to a group.

Once they have access to a group, the specified cockpit users will be able to monitor each of the resources within the group, as well as see aggregate data for the group. A cockpit user who isn't given access to any groups can log in to the cockpit, but their *My Resources* page displays zero resources. A cockpit user who is given access only to a group that contains no resources also sees a display of zero resources.

The auto-generated groups of resources (Production, Test, Development) are based on the system usage type of each resource. System usage type is configured during system installation, or later using the `global.ini` file with the `usage` parameter in the `system_information` section. Each resource belongs to an auto-generated group, however you can prevent these from being visible by configuring the cockpit settings.

Related Information

[Register a Resource \[page 70\]](#)

[Unregister a Resource \[page 76\]](#)

[Managing Resource Groups \[page 80\]](#)

[Managing Cockpit Users \[page 60\]](#)

[Configuring Cockpit Settings \[page 67\]](#)

[Configure System Usage Type \[page 304\]](#)

[Specifying Display of Auto-Generated Groups \[page 69\]](#)

5.1.5.3.1 Register a Resource

Add a resource so that cockpit users can monitor and manage it with SAP HANA cockpit.

Prerequisites

- Your cockpit user has the assigned role `Cockpit Resource Administrator` or `Cockpit Power User`. (The latter permits you to register a resource, but not assign it to a resource group).
- On the SAP HANA resource you want to register, you have created a technical user account that the cockpit will use to collect monitoring data (such as information on alerts and system performance). The technical user requires the `CATALOG READ` system privilege and `SELECT` on the `_SYS_STATISTICS` schema. SAP recommends that you set up a dedicated account for the technical user; do not allow human users to use it. See *Security Aspects of SAP HANA Cockpit*.

i Note

It's not possible by using the cockpit to create the technical user required to register a resource in the SAP HANA cockpit. You need to create this user and grant the minimum necessary authorization by using SQL as follows:

```
CREATE USER <username> PASSWORD <password> NO FORCE_FIRST_PASSWORD_CHANGE;  
GRANT CATALOG READ to <username>;  
GRANT SELECT on SCHEMA _SYS_STATISTICS to <username>
```

- If you plan to encrypt the SAP start service connection (SAP Control) or the database connection, in SAP HANA XS advanced, you have:
 1. Manually imported the server root certificate(s)
 2. Trusted the certificate(s) using the command syntax `xs trust-certificate <ALIAS> -c <CERT_FILE> [-u HTTP|JDBC]`
 3. Exported the certificate(s) to the cockpit using the commands `xs restage cockpit-hdb-service` followed by `xs restart cockpit-hdb-service`.
See *Import a Certificate for Encrypted Communication*.
- If you plan to add the resource to a group during the registration process, you have created at least one resource group.

Context

To make a resource available to cockpit users, first register the resource, then add the resource to at least one resource group, and finally assign cockpit users to the resource group.

i Note

If you are a Cockpit Power User rather than a Cockpit Resource Administrator, you will be able to register a resource only. A Cockpit Resource Administrator must add the resource to a resource group and assign cockpit users to the group.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user with the `Cockpit Resource Administrator` role or the `Cockpit Power User` role.

You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. In the Cockpit Manager, select [Registered Resources](#).
3. Select [Register Resource](#) (lower left).
4. In the Specify Resource section:
 - a. Specify whether the name of the resource should be system-generated (taking the form `database name @ system name`), or one that you will enter here.
 - b. Enter the host name for the resource you want to register.

You may need to enter a fully qualified host name. For example, if `myhost` doesn't work, try `myhost.mycompany.com`.

i Note

You can register a resource whose statistics server isn't running or is unreachable, but some cockpit features won't be available for that resource until its statistics server recovers.

- c. Choose an identifier for the resource—instance number or (for DBaaS resources) port number—and enter the identifier you specified.
 - d. Select [Single container](#) or [Multiple containers](#). If you select [Multiple containers](#):
 - o Select [System database](#) or [Tenant database](#).
 - o If you select [Tenant database](#), enter the database name.
 - e. (Optional) Enter a description of the resource.
5. In the Technical User section, enter the user name and password the cockpit will employ to collect monitoring data from the resource you're registering.
 6. In the Connection section, choose whether to encrypt the cockpit's connections to the SAP start service (SAP Control) and to the database. For information about obtaining certificates, see [Certificate Management in SAP HANA](#) in the [SAP HANA Security Guide](#).

- If you encrypt the SAP start service connection, you are allowing a secure connection (HTTPS) to the SAP start service (provided that you have met the prerequisite importing the trusted certificate(s) to the cockpit).
 - If you encrypt the database connection using a secure JDBC connection, choose whether to validate the certificate. This option lets you stipulate whether to verify that the remote server is trusted by the cockpit. Deselect the checkbox if the SAP HANA database has a certificate that differs from the one currently imported, or if you have not imported the certificate from the SAP HANA database into XS advanced. However, the recommendation is that you instead import a certificate for encrypted connections.
 - Optionally, you can enter a hostname to override the one in the certificate. You could do this to avoid the validation failure that may result from the hostname in a certificate differing from the host name that cockpit uses to connect, as in the case, for example, of a host alias, or a short hostname instead of a fully qualified domain name.
7. (Optional for Cockpit Resource Administrators) In the Groups section, select [Add to Group](#).

i Note

A resource must belong to at least one resource group—otherwise it will be invisible to cockpit users. At some point, a cockpit resource administrator will need to add this resource to a group.

8. (Optional) In the Contact section, enter the name, e-mail address, and other information for the person or group who can resolve problems and answer questions about this resource.
9. Select [Review](#).
10. Go over the details on the Register Resource Review page and use the [Edit](#) links to make changes.
11. When you're satisfied with the information on the resource, select [Register](#) (lower right).
- If the cockpit is unable to connect to the resource and the host agent reports that the resource is stopped or has an invalid or expired license, you see a confirmation request warning that the information you provided can't be validated. Click [Yes](#) to register the unreachable resource. (If you click [No](#), the cockpit does not register the resource.)
 - If the cockpit is unable to connect to the host agent to learn the status of the resource, you see a confirmation request asking for the resource's SAP HANA system ID. Enter the system ID and press [OK](#) to register the unreachable resource. (If you click [Cancel](#), the cockpit does not register the resource.)

i Note

If you enter an incorrect system ID here, or if you entered an incorrect host name, instance number, technical user, or technical user password for this resource, the cockpit allows the registration because it can't tell the difference between a host or resource that's unreachable and one that doesn't exist.

Results

When registration is successful, you can see the list of registered resources on the Resource Details page, which includes the one you just added (a fact displayed in the [Registered By](#) field. If the newly registered resource has an invalid or expired license or is offline or stopped, its software version appears as 0.00.000.00.0 (UNKNOWN).

Next Steps

- Each cockpit user accessing the resource needs to enter an existing database user name and password that has the system privilege CATALOG READ and SELECT on _SYS_STATISTICS (under *Credentials*) in order to drill down to monitor the system, unless single sign-on is in effect.
- If you've registered an offline resource, you can use the cockpit to start it. See *Start a Resource*.

Related Information

[Create a Resource Group \[page 81\]](#)

[Working with the Resource Directory \[page 165\]](#)

[Managing Cockpit Users \[page 60\]](#)

[Unregister a Resource \[page 76\]](#)

[Security Aspects of SAP HANA Cockpit \[page 86\]](#)

[Import a Certificate for Encrypted Communication \[page 94\]](#)

[Start a Resource \[page 179\]](#)

5.1.5.3.2 Edit Resource Settings, including SSO

Once a resource has been registered, you, as a cockpit resource administrator, may have reason to modify some of the original registration settings. Additionally, on a database running SAP HANA 2.0 SPS 01 you can enable or enforce single sign-on (SSO) user authentication.

Prerequisites

Your cockpit user has the assigned role `Cockpit Resource Administrator`.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user with the `Cockpit Resource Administrator` role.

You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the *Manage Cockpit* link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. On the Cockpit Manager page, click *Registered Resources*.
The Resources page lists all the systems known to the SAP HANA cockpit.

3. In the left pane, select the resource whose settings you want to modify.

i Note

If the credentials for this resource's technical user (the login the cockpit uses to collect system health and version information) need to be updated, a message on the Resource Details page alerts you to the problem. This might happen because the technical user's password has changed, for example, or the login has been deleted. If you see a message about a missing technical user or technical user password, you can't monitor the resource until you update the technical user's credentials.

4. Click *Edit*.
5. (Optional) In the Resource Details section, edit the name of the resource.
6. (Optional) In the Resource Details section, edit the description of the resource.
7. (Optional) In the Technical User section, make changes to the user name and password of the technical user, ensuring that the changes correspond to the technical user created for this resource outside of the cockpit .

i Note

It's not possible by using the cockpit to create the technical user required to register a resource in the SAP HANA cockpit. You need to create this user and grant the minimum necessary authorization by using SQL as follows:

```
CREATE USER <username> PASSWORD <password> NO FORCE_FIRST_PASSWORD_CHANGE;  
GRANT CATALOG READ to <username>;  
GRANT SELECT on SCHEMA _SYS_STATISTICS to <username>
```

8. (Optional) In the Connection section, change whether the cockpit's connections to the SAP start service (SAP Control) and to the database are encrypted. For information about obtaining certificates, see *Certificate Management in SAP HANA* in the *SAP HANA Security Guide*.
 - If you encrypt the SAP start service connection, you are allowing a secure connection (HTTPS) to the SAP start service (provided that you have met the prerequisite importing the trusted certificate(s) to the cockpit).
 - If you encrypt the database connection using a secure JDBC connection, choose whether to validate the certificate. This option lets you stipulate whether to check that the remote server is trusted by the cockpit. You can use this option if you have not imported the certificate from the SAP HANA database into XS advanced, or if the SAP HANA database has a certificate that differs from the one currently imported.
 - Optionally, you can enter a host name to override the one in the certificate. You could do this to avoid the validation failure that may result from the hostname in a certificate differing from the host name that cockpit uses to connect, as in the case, for example, of a host alias, or a short hostname instead of a fully qualified domain name.
9. (Optional) In the Contact section, edit the name, e-mail address, and other information for the person or group who can resolve problems and answer questions about this resource. You can delete the contact information when the user or users are deleted, or if the information is no longer required for any reason.
10. (Optional) In the Single Sign On section, choose *Yes* to allow SSO access to this resource rather than having cockpit prompt for a database user name and password when cockpit users are drilling down to the detailed information and monitoring capabilities in the system overview.
 - a. In the dialog box, in order to authorize this change, enter the credentials of an existing database user with the privileges:

- TRUST ADMIN
 - CERTIFICATE ADMIN
 - USER ADMIN
- b. Specify whether or not SSO should be enforced by selecting one of:

Enforce SSO Option	Description
Yes	Cockpit users must use SSO to access this resource
No	In the Resource Directory, cockpit users can choose whether to access this resource with SSO or to enter alternate database user credentials.

i Note

Enforce SSO through cockpit only after you have configured it on the database. See *Enforce Single Sign-On*.

11. Review the table listing any resource groups to which this resource currently belongs.
12. (Optional) Click [Add Resource to Group](#) to specify which resource groups should include this resource by selecting from a list of existing resource groups.
13. Click [Save](#).

Related Information

[User Authentication and Single-Sign On \[page 715\]](#)

[Enforce Single Sign-On \[page 75\]](#)

5.1.5.3.2.1 Enforce Single Sign-On

You can enforce single sign-on (SSO) user authentication on resources running SAP HANA 2.0 SPS 01 or later, provided that you follow this ordered process.

Prerequisites

- Ensure that the cockpit users who will be accessing this resource have been created.

Procedure

1. As a cockpit user with the `Cockpit Resource Administrator` role, use the Cockpit Manager to register the resource.
2. Edit the resource and enable SSO. You'll need to provide credentials for a database user with the TRUST ADMIN, USER ADMIN and CERTIFICATE ADMIN privileges.

i Note

If you enforce SSO before configuring the remote database, you will not be able to complete the setup through cockpit. (You will need to use another tool like hdbsql).

3. As any cockpit user, sign in to the cockpit.
4. Connect to the resource as an existing database user with the TRUST ADMIN, USER ADMIN and CERTIFICATE ADMIN privileges.
5. In the *Manage Users* application, set the JWT mappings.
6. Ask the newly configured database user to connect to the resource through the cockpit.
7. Return to the Cockpit Manager, edit the resource and enforce SSO (if desired).

The cockpit user will be able to connect to the resource using SSO.

Next Steps

Repeat this process for each managed resource on which you want to configure SSO.

Related Information

[Security Aspects of SAP HANA Cockpit \[page 86\]](#)

5.1.5.3 Unregister a Resource

Remove a resource from SAP HANA cockpit.

Prerequisites

Your cockpit user has the assigned role `Cockpit Resource Administrator`.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user with the `Cockpit Resource Administrator` role.

You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. In the Cockpit Manager, click [Registered Resources](#).
The Resources page lists all the systems known to the SAP HANA cockpit.
3. Click the [Unregister](#) link for the resource you're unregistering and confirm the request.
On the Resources page, you can see the list of registered systems, which no longer includes the one you just removed.

i Note

The resource will no longer be included in any resource groups, nor accessible to cockpit users.

Related Information

[Register a Resource \[page 70\]](#)

[Managing Resource Groups \[page 80\]](#)

[Working with the Resource Directory \[page 165\]](#)

5.1.5.3.4 Export Resources

Export registration information about resources in the form of a file you can use to import the resources into another system.

Prerequisites

- Your cockpit user has the assigned role `Cockpit Resource Administrator`.
- The resources you're exporting are running and available on the network.

Context

Create a `.json` export file you can import into another system running the SAP HANA cockpit..

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user with the `Cockpit Resource Administrator` role.

You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. In the Cockpit Manager, select [Registered Resources](#).
3. At the bottom of the screen, click the overflow button (the three dots to the right of [Register Resource](#)), then click [Export Resources](#).
4. (Optional) On the Export Resources screen, use the drop-down to select the group that contains the resources to be exported.
5. Select the resources you want to export and move on to the next step. Click the Resource box in the top row to select all resources in the group.
6. (Optional) Click to [Save technical user login names](#) in the export file.
Only the names of technical users are saved; you'll need to enter the technical users' passwords when you import the resources.
7. (Optional) Click to [Save contact settings](#) in the export file.
8. Click [Save export file](#).
9. At the prompt, open or save the `.json` export file.

Next Steps

Copy the `.json` export file to a location accessible to the importing system, then use the file to import the resources to the new system.

Related Information

[Import Resources \[page 79\]](#)

5.1.5.3.5 Import Resources

Add multiple resources so that users can monitor and manage them with SAP HANA cockpit.

Prerequisites

- Your cockpit user has the assigned role `Cockpit Resource Administrator`.
- You have created a `.json` export file on another SAP HANA system that specifies the resources to be imported, and you have copied that file to a location accessible to the importing system.
- If you plan to encrypt the SAP start service connection (SAP Control) or the database connection, in SAP HANA XS advanced, you have:
 1. Manually imported the server root certificate(s)
 2. Trusted the certificate(s) using the command syntax `xs trust-certificate <ALIAS> -c <CERT_FILE> [-u HTTP|JDBC]`
 3. Exported the certificate(s) to the cockpit using the commands `xs restage cockpit-hdb-service` followed by `xs restart cockpit-hdb-service`.See [Import a Certificate for Encrypted Communication \[page 94\]](#).
- The resources you're importing are running and available on the network.

Context

Adding resources to groups enables users to access them. To make a resource available to cockpit users, first register the resource, then add the resource to at least one resource group, and finally assign cockpit users to the resource group.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user with the `Cockpit Resource Administrator` role.

You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```
2. In the Cockpit Manager, select [Registered Resources](#).
3. At the bottom of the screen, click the overflow button (the three dots to the right of [Register Resource](#)), then click [Import Resources](#).
4. On the Import File screen, browse to the `.json` file that specifies the resources to be imported. Click [Step 2](#) to continue.

The cockpit displays a list of resources from the file.

5. Select the resources you want to register and click the next step.
6. For each resource you're importing, enter the technical user name (if prompted) and password.
7. (Optional) Enter the name and contact information of someone responsible for each resource you're importing.
8. Click [Review](#) to check what you've entered.

To add or edit optional information like the description and contact details, click [Edit](#) at the bottom of the page.

9. Click [Import Resources](#) (bottom right) to register the imported resources.

Next Steps

(Optional) Add the imported resources to groups.

Related Information

[Export Resources \[page 77\]](#)

[Managing Resource Groups \[page 80\]](#)

5.1.5.3.6 Managing Resource Groups

As a cockpit user with the `Cockpit Resource Administrator` role, you can create, populate, or remove the groups used to grant resources access to other cockpit users.

A resource group—a named set of one or more registered resources—controls management and monitoring privileges. When you assign a cockpit user to a resource group that you have created, you enable the user to monitor and manage the group's resources through the cockpit.

i Note

A cockpit user has no access to a resource unless the user and the resource belong to the same resource group.

Assigning resources and cockpit users to resource groups let you, as a cockpit resource administrator:

- View and administer similar resources together.
- Control which cockpit users can see and use particular resources.

You can create resource groups whatever way you want—for example, by groups based on resources' geographic location, ownership, or purpose.

Each resource also belongs to an auto-generated group, however you can prevent these from being visible by configuring the cockpit settings. These auto-generated groups of resources (Production, Test, Development) are based on the system usage type of each resource. System usage type is configured during system

installation, or later using the `global.ini` file with the `usage` parameter in the `system_information` section.

i Note

In order to have access to a resource, a cockpit user must be granted access to a resource group that you create and to which you assign the resource. You cannot assign cockpit users to an auto-generated group, or to individual resources.

Related Information

[Create a Resource Group \[page 81\]](#)

[Add or Remove Resources in Resource Groups \[page 82\]](#)

[Delete a Resource Group \[page 83\]](#)

[Add or Remove Users in Resource Groups \[page 84\]](#)

[Managing Cockpit Users \[page 60\]](#)

[Specifying Display of Auto-Generated Groups \[page 69\]](#)

5.1.5.3.6.1 Create a Resource Group

Set up a group you can use to display, manage, and control access to related resources.

Prerequisites

Your cockpit user has the assigned role `Cockpit Resource Administrator`.

Context

When you create a resource group, you can add both resources and cockpit users. Only users assigned to the group can see and access the group's resources.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user with the `Cockpit Resource Administrator` role.

You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. In the Cockpit Manager, select [Resource Groups](#).
3. Click [Create Group](#) (bottom of page).
4. On the Create Group page, enter a name in the Name field.
In a group name, you can use uppercase and lowercase letters, the digits 0 through 9, underscores, hyphens, periods, and spaces.
5. (Optional) Enter a description of the group in the Description field.
6. Click to expand the next step.
7. (Optional) Click [Add Resources](#) and select registered resources to add to the new group.
8. Click to expand the next step.
9. (Optional) Click [Add Users](#) and select cockpit users who will have access to the new group.
10. Scroll up and review what you entered.
11. Click [Create Group](#).

The cockpit displays the new group on the Resource Groups page.

Related Information

[Add or Remove Resources in Resource Groups \[page 82\]](#)

[Add or Remove Users in Resource Groups \[page 84\]](#)

[Remove a Resource Group \[page 83\]](#)

5.1.5.3.6.2 Add or Remove Resources in Resource Groups

Add a resource to a resource group or remove a resource from a resource group.

Prerequisites

Your cockpit user has the assigned role `Cockpit Resource Administrator`.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user with the `Cockpit Resource Administrator` role.

You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. In the Cockpit Manager, select [Resource Groups](#).
3. In the left pane, select the group you want to modify.

Only the resource groups created by administrators appear on the Resource Groups page. You cannot modify or delete the auto-generated groups based on configured system usage type (Production, Development, Test), nor the My Resources group, which is a compilation of all groups that a user has access to.

4. (Optional) To add a resource:
 - a. If necessary, click [Resources](#) to display the list of resources in this group.
 - b. Click [Add Resource](#) (far right).
 - c. Select the resources to add and click [OK](#).

The cockpit displays the group, including the newly added resource or resources, on the Resource Groups screen.

5. (Optional) To remove a resource:
 - a. If necessary, click [Resources](#) to display the list of resources in this group.
 - b. Click the red  to the right of the resource to be deleted and confirm the deletion.

The cockpit displays the group, without the removed resource, on the Resource Groups screen.

Related Information

[Create a Resource Group \[page 81\]](#)

[Remove a Resource Group \[page 83\]](#)

[Working with the Resource Directory \[page 165\]](#)

[Search, Sort, and Filter Tools for Resources and Groups \[page 167\]](#)

5.1.5.3.6.3 Delete a Resource Group

Delete a resource group from the SAP HANA cockpit.

Prerequisites

Your cockpit user has the assigned role `Cockpit Resource Administrator`.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user with the `Cockpit Resource Administrator` role.

You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. In the Cockpit Manager, select [Resource Groups](#).
3. In the left pane, select the group you want to remove.
4. Click [Delete](#) (upper right) and confirm the deletion.
On the Resource Groups page, you can see the list of groups, which no longer includes the one you just removed. The registered resources from the group will be available for you to include in other groups.

Note

Any cockpit users who had access to these resources solely through this resource group will not be able to access them.

Related Information

[Create a Resource Group \[page 81\]](#)

[Add or Remove Resources in Resource Groups \[page 82\]](#)

[Working with the Resource Directory \[page 165\]](#)

[Search, Sort, and Filter Tools for Resources and Groups \[page 167\]](#)

5.1.5.3.6.4 Add or Remove Users in Resource Groups

Add a cockpit user to a resource group or remove a user from a resource group.

Prerequisites

- Your cockpit user has the assigned role `Cockpit Resource Administrator`.
- The `Cockpit User Administrator` has created some cockpit users. See [Managing Cockpit Users](#).

Context

Only cockpit users who are assigned to a resource group can see the group's registered resources in the cockpit.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit user with the `Cockpit Resource Administrator` role.

You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation, or by following the [Manage Cockpit](#) link in the cockpit. The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

2. In the Cockpit Manager, select [Resource Groups](#).
3. In the left pane, select the group you want to modify.

Only the resource groups created by administrators appear on the Resource Groups page. You cannot modify or delete the auto-generated groups based on configured system usage type (Production, Development, Test), nor the My Resources group, which is a compilation of all groups that a user has access to.

4. (Optional) To add a cockpit user:
 - a. If necessary, click [Users](#) to display the list of cockpit users in this group.
 - b. Click [Add User](#) (far right).
 - c. Select the users to add and click [OK](#).

The cockpit displays the group, including the newly added cockpit user(s), on the Resource Groups page.

5. (Optional) To remove a user:
 - a. If necessary, click [Users](#) to display the list of cockpit users in this group.
 - b. Click the red ⊗ to the right of the user to be deleted and confirm the deletion.

The cockpit displays the group, without the removed user, on the Resource Groups page.

Related Information

[Managing Resource Groups \[page 80\]](#)

[Managing Cockpit Users \[page 60\]](#)

5.1.6 Security Aspects of SAP HANA Cockpit

Security considerations for SAP HANA cockpit include user management, single sign-on and certificate management.

User Authentication

Several types of credentials are used within SAP HANA cockpit:

Credential	Details
COCKPIT_ADMIN	The master user for the cockpit created during the installation process. The password for the cockpit administrator user is the master password established during the installation process. This master user is assigned all three administrator roles, and can therefore access all aspects of the Cockpit Manager, and can create users, register resources, and assign users and resources to resource groups.
Cockpit Users	The business users with access to the cockpit. Each is assigned one or more roles: <ul style="list-style-type: none">• The cockpit administrator role can access the Cockpit Settings section of the Cockpit Manager, where they can configure the cockpit.• The cockpit resource administrator role can access the Registered Resource and Resource Groups sections of the Cockpit Manager, where they can register resources, create resource groups, and assign cockpit users to resource groups.• The cockpit user administrator role can access the Manage User section of the Cockpit Manager, where they can create and manage cockpit users.• The cockpit user role can access the SAP HANA cockpit, where they can view the resources in the resource groups to which they have been granted access.• The cockpit power user role can access the SAP HANA cockpit and the Registered Resource section of the Cockpit Manager

Credential	Details
Technical User	<p>An application global per-resource set of database credentials for access to remote resources and necessary to regularly gather information such as state, status, and other generalized KPIs.</p> <p>The technical user is created outside of the cockpit, but the technical user must be specified when a resource is registered in the cockpit. The technical user requires the privileges necessary to perform registration and statistics gathering: at a minimum, CATALOG READ system privilege and SELECT permission on _SYS_STATISTICS catalog.</p>
Database User (User Remote Login)	<p>The per resource/per user set of database credentials for a remote resource used by the cockpit user to view more sensitive information, and to make changes within their roles as defined on that resource.</p> <p>Each cockpit user needs to provide database user credentials with the system privilege CATALOG READ and SELECT on _SYS_STATISTICS in order to drill down in the cockpit to the overview information for a specific resource. The cockpit securely encrypts and stores separate database credentials for each cockpit user, but you can clear and re-enter the credentials through the <i>Resource Directory</i> when you desire to do so.</p>
Single Sign-On (SSO)	<p>For any systems running version SAP HANA 2.0 SPS 01, or later, the cockpit resource administrator has the option to enable or enforce SSO. See <i>Setting Up Single Sign-On</i>.</p>
Operating System User	<p>A per resource set of credentials for accessing the SAP Control process (starting and stopping the resource, and restoring features). This is usually the <sid>adm account. The cockpit securely encrypts and stores these credentials, but you can clear and re-enter the credentials through the <i>Resource Directory</i> when you desire to do so.</p>
Internal Communication	Service-to-service authentication
SAP HANA Service Broker User	For application persistence using the application's SAP HANA express database

Note

It's not possible by using the cockpit to create the technical user required to register a resource in the SAP HANA cockpit. You need to create this user and grant the minimum necessary authorization by using SQL as follows:

```
CREATE USER <username> PASSWORD <password> NO FORCE_FIRST_PASSWORD_CHANGE;
GRANT CATALOG READ to <username>;
GRANT SELECT on SCHEMA _SYS_STATISTICS to <username>
```

Network and Communication Security

The cockpit uses secure protocols on all client browser connections to HTTPS ports. Communication to SAP HANA databases uses JDBC, and may be secured by importing certificates into the cockpit. Additional communication is made to the remote hosts using a restful interface which also may be secured. You can use properly signed certificates for the cockpit's external ports as well. For more information about obtaining certificates, see *Certificate Management in SAP HANA* in the *SAP HANA Security Guide*.

In a large enterprise it's likely that you may generate internal certificates that are signed by an internal certificate signing authority. In this case, you could insert the single (root) certificate from the signing authority. Any certificates signed by that authority (such as HTTPS or JDBC certificates) are automatically trusted. However, in a default installation, the SAP HANA system generates a self-signed certificate. In this situation, if the certificate is not replaced by a correctly signed one, then that specific certificate should be imported in order to enable trust.

Related Information

[Managing Cockpit Users \[page 60\]](#)

[Edit Settings for a Cockpit User \[page 64\]](#)

[Working with the Resource Directory \[page 165\]](#)

[Import a Certificate for Encrypted Communication \[page 94\]](#)

[Setting Up Single Sign-On \[page 88\]](#)

5.1.6.1 Setting Up Single Sign-On

Single sign-on (SSO) is a form of authentication which allows user access without requiring that the user enter credentials every time.

The SAP HANA cockpit offers the option to configure SSO to access cockpit itself (including the cockpit manager), and to connect to a registered resource.

Configuring SSO access to the cockpit itself means that you need not provide cockpit user credentials in order to access the cockpit or the cockpit manager.

The option to enable or enforce single-sign on (SSO) for a specific resource removes the need for providing database user credentials each time you connect to the resource. If you enforce SSO, cockpit users must use SSO to access the resource. If you enable SSO, but do not enforce it, cockpit users can choose whether to access this resource with SSO or to enter alternate database user credentials.

Related Information

[Edit Resource Settings, including SSO \[page 73\]](#)

[Enforce Single Sign-On \[page 75\]](#)

[Configure SSO Access to the SAP HANA Cockpit \[page 89\]](#)

[Configure SSO Access to a Resource \[page 92\]](#)

5.1.6.1.1 Configure SSO Access to the SAP HANA Cockpit

To configure single-sign on authentication to the SAP HANA cockpit, use the cockpit for SAP HANA extended application services, advanced model (XS advanced cockpit) and the SAP Cloud Platform Identity Authentication Administration Console. These are external tools and not part of the SAP HANA cockpit itself.

Context

Authentication standard supported for SSO to SAP HANA cockpit: SAML

For more information about the XS advanced cockpit, including standards supported for SSO, refer to *Maintaining the XS Advanced Runtime Environment with SAP HANA XS Advanced Cockpit* or *Managing SAML Identity Providers in XS Advanced* in the *SAP HANA Administration Guide*.

i Note

The steps in these instructions detail how to configure SSO access using the SAP Cloud Platform Identity Authentication Administration Console. If you are using another identity provider (IDP), the steps or details may vary.

Procedure

1. Log in to the SAP Cloud Platform Identity Authentication Administration Console and create named groups.
 - a. Navigate to
`https://<IDP_URL>/admin`
 - b. Select *User Groups*.
 - c. Use *Add +* to provide a name for each group. For example:
 - HANA_COCKPIT_ADMIN
 - HANA_COCKPIT_RESOURCE_ADMIN
 - HANA_COCKPIT_USER_ADMIN
 - HANA_COCKPIT_USER
 - HANA_COCKPIT_POWER_USER

In the example above we use cockpit roles for group names; you can choose your own group names.
 - d. Select *Save*.
 - e. Exit the SAP Cloud Platform Identity Authentication Administration Console.
2. Use the XS advanced cockpit to add a new SAML identity provider.

- a. Retrieve SAML2 metadata from your IDP.

For example, using the SAP IDP, navigate to

```
https://<IDP_URL>/saml2/metadata
```

- b. Use the `xsa` command to access the xsa-cockpit URL.
 - c. In the left pane, select **Security > Trust Configuration**.
 - d. Select **New Trust Configuration** to add a new provider.
 - e. In the metadata text box, enter XML based on the SAML metadata you retrieved from the IDP in step 2.a [page 90].
 - f. Select **Parse** to fill in the remaining fields.
 - g. Select **Save**.
3. In the XS advanced cockpit, map the role collections. (For more information, see SAP Note 2569903).
 - a. Select **Security > Trust Configuration**.
 - b. Select the trust configuration you created in step 2 [page 89].
 - c. Select **Role Collection Mappings**.
 - d. Select **New Role Collection Mapping** to map each of the role collections to a group within the IDP. If you use the sample group names in step 1.c [page 89], the mappings are:

Role Collection	Group Name
COCKPIT_ADMIN	HANA_COCKPIT_ADMIN
COCKPIT_RESOURCE_ADMIN	HANA_COCKPIT_RESOURCE_ADMIN
COCKPIT_USER_ADMIN, XS_USER_ADMIN	HANA_COCKPIT_USER_ADMIN
COCKPIT_USER	HANA_COCKPIT_USER
COCKPIT_POWER_USER	HANA_COCKPIT_POWER_USER

Note

If you're using the COCKPIT_USER_ADMIN role collection, you must map both COCKPIT_USER_ADMIN and XS_USER_ADMIN to the IDP group representing the cockpit user administrator, in this case HANA_COCKPIT_USER_ADMIN.

- e. Select **Save**.
4. Retrieve spring SAML metadata from


```
https://<cockpit_FQDN><instance#>32/uaa-security/saml/metadata
```

For example, if the cockpit is running on yourserver.company.com instance 03, the URL is

```
https://yourserver.company.com:30332/uaa-security/saml/metadata
```
 5. The file for use in the IDP configuration downloads automatically.
 6. Log in to the SAP Cloud Platform Identity Authentication Administration Console and create a new application to represent the SAP HANA cockpit.
 - a. Select **Applications and Resources**.
 - b. Select **Add Application**.
 - c. Enter an application name.
 - d. Select **Save**.

7. In the SAP Cloud Platform Identity Authentication Administration Console, configure SAML.
 - a. Select the new application.
 - b. Select *SAML 2.0 Configuration*.
 - c. Select the browse button to locate and select the downloaded metadata file (~/Downloads/spring_saml_metadata.xml).

i Note

Only one application may be configured with the SAML provider. An error will occur if you introduce a duplicate, even with a different name.

8. In the SAP Cloud Platform Identity Authentication Administration Console, configure the application Name ID attribute.
 - a. Select *Name ID Attribute*.
 - b. Select *Email*.
 - c. Select *Save*.
9. In the SAP Cloud Platform Identity Authentication Administration Console, provide assertion attributes.
 - a. Select *Assertion Attributes*.
 - b. Add a groups attribute.
 - c. Modify the *groups* attribute (lowercase) to Groups (title case).
 - d. Accept the default assertion attribute for *First name*, *Last name* and *E-mail*.
10. In the SAP Cloud Platform Identity Authentication Administration Console, add the user.
 - a. Select *User Management*.
 - b. Select *Add User*.
 - c. Provide last name, email and user type, and account activation options.
 - d. Press *Save*.
11. In the SAP Cloud Platform Identity Authentication Administration Console, edit the user.
 - a. Select *Applications*.
 - b. Add the application you created in the XSA Admin tools.
 - c. Select *User Groups* and add the groups appropriate for this user (corresponding to the cockpit role).
12. Log in to the SAP HANA cockpit. You see the link on the sign-in page, and do not need to enter cockpit user credentials. (Selecting the link might bring up the IDP authentication page.)

Related Information

[Setting Up Single Sign-On \[page 88\]](#)

[Configure SSO Access to a Resource \[page 92\]](#)

[SAP Note 2569903](#) 

5.1.6.1.2 Configure SSO Access to a Resource

Enabling single sign-on (SSO) allows a cockpit user to log on to a resource without being prompted for database user credentials.

Prerequisites

- The resource has been registered in the cockpit. It meets these version restrictions:
 - SAP HANA 1.0 SPS 12 revision 14 or later, or
 - SAP HANA 2.0 SPS 01 or later
- You have a database user with the CATALOG READ, TRUST ADMIN, CERTIFICATE ADMIN, and USER ADMIN privileges.

The database user we use in the steps below is SSO_USER.

To assign the necessary system privileges to an existing user called SSO_USER, execute these SQL statements:

```
GRANT TRUST ADMIN TO SSO_USER;  
GRANT CERTIFICATE ADMIN TO SSO_USER;  
GRANT USER ADMIN TO SSO_USER;  
GRANT CATALOG READ TO SSO_USER;
```

- You have a second database user with the USER ADMIN system privilege.
The second database user we use in the steps below is USER_ADMIN.
- A user with the Cockpit Administrator or Cockpit User Administrator role has created one or more cockpit users who will be accessing this resource and assigned them the role `Cockpit User`.
The cockpit user we use in the steps below is COCKPIT_USER.
- For step 1 [page 93], you have a cockpit user with the Cockpit Administrator or Cockpit Resource Administrator role.
The cockpit administrator we use in the steps below is COCKPIT_ADMIN.

Context

Authentication standard supported for SSO to a managed resource: JSON Web Token (JWT).

i Note

Before enabling SSO, consider migrating the Personal Security Environment (PSE) file to an in-database store. When SSO is enabled, a new PSE file may be created, which may prevent cockpit access to stored certificates. See SAP Note 265666.

The steps below configure SSO for a single managed SAP HANA resource. You'll need to perform them for each resource where SSO is needed. There's a separate set of steps for configuring SSO for the cockpit itself—see the Related Links at the end.

The result of this process is to link the cockpit account we're calling COCKPIT_USER with the database account we're calling SSO_USER. When SSO is enabled, you'll be able to sign in to the cockpit as COCKPIT_USER by entering the password, and from there sign in to a resource as SSO_USER without providing credentials. You

can also sign in to the resource using different database credentials. If you choose to enforce SSO, you will be able to sign in to the resource from the cockpit **only** as SSO_USER.

Procedure

1. Enable SSO in the Cockpit Manager:
 - a. Log in to the Cockpit Manager as COCKPIT_ADMIN.
 - b. Select *Registered Resources*.
 - c. In the left column, choose the resource for which you want to enable SSO.
 - d. Click *Edit* (bottom of screen).
 - e. Scroll down to the Single Sign On section and select *Enable SSO: Yes*.
 - f. In the Authorize SSO Setting Change dialog, enter the credentials of a cockpit user that has the CATALOG READ, TRUST ADMIN, CERTIFICATE ADMIN, and USER ADMIN privileges. We're using SSO_USER here, as described in the Prerequisites.
 - g. Click *Save*.
2. Sign in to the database as a user with the USER ADMIN system privilege—we're using USER_ADMIN here.
3. On the System Overview page for the resource, click *Manage users* (under User & Role Management).
4. Select the user with the CATALOG READ, TRUST ADMIN, CERTIFICATE ADMIN, and USER ADMIN privileges—SSO_USER for purposes of this procedure.
5. To set the JWT mappings:
 - a. Under Authentication, select *JWT - You must add at least one identity provider*.
 - b. Click *Add JWT Identity*.
 - c. Select XS_APPLICATIONUSER from the *Identity Provider* dropdown.
 - d. Turn off the *Automatic Mapping by Provider*.
 - e. Enter the user name of an existing cockpit user in *External Identity*—we're using COCKPIT_USER here—and click *Save*.
6. Log out of the cockpit and log in as COCKPIT_USER (the cockpit user for which you just set up JWT).
7. Go to the Resource Directory and click the *Choose Authentication* link for the resource you're configuring.
8. Make sure *Log on via single sign on* is selected and click *OK*.
9. Click the resource name to log in.

The resource signs you in as SSO_USER, though you logged in to the cockpit as COCKPIT_USER.
10. (Optional) Enforce SSO through the Cockpit Manager:
 - a. On the Home page, under Manage Landscape, click *Manage Cockpit* to return to the Cockpit Manager.
 - b. Click *Registered Resources*.
 - c. Choose the resource for which you want to enforce SSO and click *Edit* (lower right).
 - d. Under Single Sign On, select *Enforce SSO: Yes* and click *Save*.
 - e. In the Authorize SSO Setting Change dialog, enter the credentials of SSO_USER.
 - f. Click *Go to SAP HANA Cockpit* (lower right) to return to the cockpit.
 - g. Open the Resource Directory and find the resource for which you've configured SSO.

The resource's Credentials column now says `SSO enforced`. You can access the database only as COCKPIT_USER's mapped database user, SSO_USER—the cockpit does not allow you to enter other credentials.

Next Steps

Repeat these steps as needed to enable SSO for other resources.

Related Information

[Create or Enable a Cockpit User \[page 62\]](#)

[Configure SSO Access to the SAP HANA Cockpit \[page 89\]](#)

[SAP Note 2656666](#) 

5.1.6.2 Import a Certificate for Encrypted Communication

You can import a certificate to enable encrypted communication from the cockpit to the remote SAP start service (SAP Control), or communication from the cockpit to an SAP HANA database

Procedure

1. Sign on as `<sid>adm` to the remote system with which you want the cockpit to establish a connection.
2. On the remote system, run the `sapgenpse` tool to export the certificate(s) from the in-memory certificate collection, or from the file-system PSE (`$SECUDIR/sapgenpse export_own_cert`).
The location of the required certificate(s) depends on how you manage certificates in your system. For instance, you may be able to open a browser, point to SAP Control, and obtain the certificate.
3. Sign on as `<sid>adm` to the system hosting the cockpit.
4. Run the command `xs login`, and provide the name and password of the `COCKPIT_ADMIN` so that the `<sid>adm` can execute the `xs` command line tool in the user context of `COCKPIT_ADMIN`.
5. In SAP HANA XS advanced model (XSA), trust the certificates using the command syntax `xs trust-certificate <ALIAS> -c <CERT_FILE> [-u HTTP|JDBC]`, where `<ALIAS>` is a unique alias for the certificate within XSA, for example. `BZ1_SAPCONTROL` and `<CERT_FILE>` is the certificate you exported from the remote system.
6. Run the commands `xs restage cockpit-hdb-service cockpit-hdb-svc` followed by `xs restart service cockpit-hdb-svc` to export the certificates to the cockpit.

See *XS CLI: Certificates* in the *SAP HANA Developer Guide (For SAP HANA XS Advanced Model)*. For information about obtaining certificates, see *Certificate Management in SAP HANA* in the *SAP HANA Security Guide*.

Related Information

[Register a Resource \[page 70\]](#)

5.1.6.3 Ensuring a Secure Browser Connection

After you've installed the SAP HANA cockpit, ensure that communication with your web browser is encrypted.

Context

Unencrypted communication with a browser could pose a security risk. To secure the connection, provide the XSA server with a certificate that is signed by a certificate authority, or use your own self-signed certificate.

Procedure

1. To use a certificate from a certificate authority:
 - a. Obtain a signed certificate from the certificate authority.
 - b. In SAP HANA XS advanced model (XSA), trust the certificate using the command syntax `xs trust-certificate <ALIAS> -c <CERT_FILE> -u HTTP` where <ALIAS> is a unique alias for the certificate within XSA, and <CERT_FILE> is the certificate you obtained.
2. To use your own certificate:
 - a. Create a self-signed root certificate.
 - b. Create a second certificate that is signed by the root certificate
 - c. In SAP HANA XS advanced model (XSA), trust the certificate using the command syntax `xs trust-certificate <ALIAS> -c <CERT_FILE> -u HTTP` where <ALIAS> is a unique alias for the certificate within XSA, and <CERT_FILE> is the second signed certificate.
 - d. Update your browser's list of trusted root certificates to include the newly-created root certificate file.

5.1.6.4 Data Protection and Privacy in SAP HANA Cockpit

SAP HANA cockpit provides tools you can use to conform to legal and business requirements for protecting personal data stored in the system.

Introduction

Data protection is associated with numerous legal requirements and privacy concerns. In addition to compliance with general data privacy regulation, it is necessary to consider compliance with industry-specific

legislation in different countries. SAP provides specific features and functions to support compliance with regard to relevant legal requirements, including data protection. SAP does not give any advice on whether these features and functions are the best method to support company, industry, regional, or country-specific requirements. Furthermore, this information should not be taken as advice or a recommendation in regards to additional features that would be required in specific IT environments; decisions related to data protection must be made on a case-by-case basis, taking into consideration the given system landscape and the applicable legal requirements.

i Note

SAP does not provide legal advice in any form. SAP software supports data protection compliance by providing security features and specific data protection-relevant functions, such as simplified blocking and deletion of personal data. In many cases, compliance with applicable data protection and privacy laws will not be covered by a product feature. Definitions and other terms used in this document are not taken from a particular legal source.

Glossary

Term	Definition
Consent	The action of the data subject confirming that the usage of his or her personal data shall be allowed for a given purpose. A consent functionality allows the storage of a consent record in relation to a specific purpose and shows if a data subject has granted, withdrawn, or denied consent.
Deletion	The irreversible destruction of personal data.
Personal data	Any information relating to an identified or identifiable natural person ("data subject"). An identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural, or social identity of that natural person.
Purpose	A legal, contractual, or in other form justified reason for the processing of personal data. The assumption is that any purpose has an end that is usually already defined when the purpose starts.

User Consent

SAP HANA cockpit stores only personal data entered by users; it never collects personal data without a user's knowledge.

Logging Read Access and Changes

SAP HANA cockpit provides tools for auditing access and changes to personal data stored in SAP HANA databases. For details, see *Auditing Database Activity*.

Information Report

SAP HANA cockpit does not store any personal data except what is entered (if anything) in the optional contact information for resources. This typically includes the name and e-mail address of the contact person, and you can see it in the resource registration.

Deletion of Personal Data

You can remove unneeded user accounts and resource contact information. See:

- [Edit Resource Settings, Including SSO](#)
- [Delete a Cockpit User or Revoke Cockpit Access](#)

Related Information

[Edit Resource Settings, including SSO \[page 73\]](#)

[Delete a Cockpit User or Revoke Cockpit Access \[page 65\]](#)

5.1.6.5 Auditing in SAP HANA Cockpit

Audit logging lets you track events like logins and creation and deletion of user accounts.

Prerequisites

- You have the password for the `COCKPIT_ADMIN` user, or
- If you want to use the Audit Log as a user other than `COCKPIT_ADMIN`, you have created an appropriate role collection and assigned it to that user. Give the role collection the application role `AuditLogViewer`, which is part of application `auditlog-ui` and role template `AuditLogViewer`. For details on building role collections, see *Maintain Roles for XS Advanced Applications* in the *SAP HANA Administration Guide*.

Context

You can use the Audit Log to identify log entries for events that you want to track.

Procedure

1. Log in to SAP HANA express edition or XS advanced:

```
xs login [-a <API_URL>] [-u <username>] [-p <password>] [-o <organization>] [-s <space>]
```

where `<API_URL>` is the API endpoint (for example, `https://api.example.com`)

2. To list applications, enter:

```
xs a
```

3. In the results, find the entry for the auditlog-ui service, which includes a URL.
4. Enter the auditlog-ui URL in a browser to open the *Audit Log*.
5. Use the Audit Log sorting and filtering tools to find logged events of interest to you.

5.1.7 Managing Resources, Users, and Groups with the Cockpit APIs

SAP HANA cockpit provides modifying (POST) and nonmodifying (GET) REST APIs. You access the APIs differently depending on whether you're calling a POST or a GET API and whether you're calling it programmatically or from a browser.

Using Cockpit APIs from an External Program

The external tool you use to invoke SAP HANA cockpit APIs must be capable of sending complex REST calls. The tool might be a custom program or a browser with a REST console plug-in. A browser with a REST plug-in will be useful for testing.

To access cockpit POST and GET APIs from an external program, you must:

1. Obtain a service key for the external tool. (You do this only once.)
2. Obtain an OAuth token using information in the service key, plus a cockpit user with the appropriate role. (You can use the OAuth token until it expires—usually in 30 minutes.)
3. Using the OAuth token, invoke one or more cockpit APIs.

Obtaining a Service Key

The cockpit runs on SAP HANA extended application services, advanced model (XS advanced), which provides authentication and authorization to external tools via service keys to the User Authentication and Authorization component (cockpit-UAA). You use `xs` commands in SAP HANA XS advanced to generate a service key for your external tool. The service key is a block of Javascript Object Notation (JSON) code; you store it in a file that your application can load and use.

Run the `xs create-service-key cockpit-uaa <tool-name>-cockpit-uaa` command to generate a service key:

```
xs create-service-key cockpit-uaa <tool-name>-cockpit-uaa
```

Next, store the service key in a `.json` file:

```
xs service-key cockpit-uaa <tool-name>-cockpit-uaa > <tool-name>-service-key.json
```

Finally, edit the `.json` file to remove the lines before and after the curly brackets. The original file takes this form:

```
Getting service key "foo-cockpit-uaa" for service instance "cockpit-uaa" ...
{
  "tenantmode" : "dedicated",
  "clientid" : "sb-cockpit!il",
  "verificationkey" : "<REDACTED>",
  "xsappname" : "cockpit!il",
  "identityzone" : "uaa",
  "identityzoneid" : "uaa",
  "clientsecret" : "<REDACTED>",
  "url" : https://host:30032/uaa-security
}
```

Change it to look similar to this:

```
{
  "tenantmode" : "dedicated",
  "clientid" : "sb-cockpit!il",
  "verificationkey" : "<REDACTED>",
  "xsappname" : "cockpit!il",
  "identityzone" : "uaa",
  "identityzoneid" : "uaa",
  "clientsecret" : "<REDACTED>",
  "url" : https://host:30032/uaa-security
}
```

Obtaining an OAuth Token

OAuth tokens are granted to cockpit users, so you must identify a cockpit user you can employ to run the cockpit APIs. To find the role needed for the API you want to run, see [Using Cockpit POST APIs \[page 100\]](#) or [Using Cockpit GET APIs \[page 103\]](#).

Using your external tool along with the UAA URL provided in the service key and the cockpit user you've identified, invoke the UAA API `POST /oauth/token` with a `grant_type` of `password`. The `clientid` and `clientsecret` come from the service key (above). For example:

```
POST https://host:30032/uaa-security/oauth/token HTTP/1.1
Host: localhost:8080
Accept: application/json
Authorization: Basic YXBwOmFwcGNsaWVudHhNlY3JldA==
"grant_type=password&username=marissa&password=koala&clientid=<clientid>
&clientsecret=<clientsecret>"
```

Expect a response code of 200 with a body similar to this:

```
{
  "access_token": "2YotnFZFEjr1zCsicMWpAA",
```

```
"token_type": "bearer",
"expires_in": 3600
}
```

→ Tip

Notice the value of `expires_in`, which tells you how long (in seconds) you can use the token before getting a new one.

For more on the `POST /oauth/token` API, see <https://github.com/cloudfoundry/uaa/blob/master/docs/UAA-APIs.rst>.

Invoking Cockpit APIs

To obtain the URL for the `cockpit-adminui-svc`, which you'll use to access the cockpit APIs programmatically as described above, issue this command to SAP HANA XS advanced:

```
$ xs apps | grep cockpit-adminui-svc
cockpit-adminui-svc          STARTED          1/1          128 MB
<unlimited> https://host:51025
```

Related Information

[Using Cockpit POST APIs \[page 100\]](#)

[Using Cockpit GET APIs \[page 103\]](#)

5.1.7.1 Using Cockpit POST APIs

Details on the SAP HANA cockpit APIs that create, delete, or change objects in the cockpit..

For instructions on setting up programmatic access to the cockpit APIs, see [Managing Resources, Users, and Groups with the Cockpit APIs \[page 98\]](#).

The APIs described in the table below accept HTTP POST operations with arguments passed in the body in JSON format. Each API returns data in JSON format.

Success

If an API succeeds, it returns:

- Status 200
- A JSON response in this form:

Sample Code

```
{
  result: {
    <some data>
  }
}
```

Failure

If an API fails, it returns:

- One of these statuses:
 - 401 – Unauthorized (the token provided is not accepted)
 - 403 – Forbidden (the token does not include appropriate scopes to execute this API)
 - 400 – Bad request (an input argument is missing or invalid)
 - 500 – Server error (something else went wrong)
- A JSON response in this form:

Sample Code

```
{
  message: {
    <error message with properties such as resource key, default text,
    etc.>
  }
}
```

SAP HANA Cockpit POST APIs

Each API listed in the POST APIs table accepts HTTP POST operations with arguments passed in the body in JSON format. Each API returns data in JSON format. These POST APIs require you to use an authentication token as described in [Managing Resources, Users, and Groups with the Cockpit APIs \[page 98\]](#). Because out-of-the-box Web browsers don't support POST operations, you need a suitable plug-in to use or test the cockpit's POST APIs with a browser.

POST APIs

API	What It Does	Input Parameters	Response ¹	Required Role
/registration/ SystemRegister	Registers an SAP HANA resource with the cockpit using JDBC. Post to the API with parameters in the body in JSON format (Content-Type=application/json; charset=UTF-8).	<p>hostName</p> <p>instanceNumber or port</p> <p>techUser</p> <p>techUserCredentials</p> <p>isMultiTenant – can be omitted if you specify port instead of instanceNumber</p> <p>databaseName – can be omitted if you specify port instead of instanceNumber</p> <p>security (encryptJDBC, validateServerCertificate, hostNameInCertificate)</p>	resid – resource ID of newly registered resource	Cockpit Resource Administrator or Cockpit Power User
/registration/ SystemUnregister	Unregisters an SAP HANA resource. Post to the API with parameters in the body in JSON format (Content-Type=application/json; charset=UTF-8).	resid – integer, required. The unique resource ID previously returned by a registration API.		Cockpit Resource Administrator or, if you registered the resource, Cockpit Power User
/registration/ ResourceUnregister	Unregisters an SAP HANA resource. Deprecated—use SystemUnregister instead.	resid		Cockpit Resource Administrator
/group/ GroupCreate	Creates a new resource group	<p>groupName</p> <p>groupDescription</p>	groupId	Cockpit Resource Administrator
/group/ GroupDelete	Deletes a resource group	groupId		Cockpit Resource Administrator

¹ See also the general responses described above under Success and Failure.

API	What It Does	Input Parameters	Response ¹	Required Role
/user/ CockpitUserCreate	Creates a new cockpit user	username password email roleCollections []	cockpitId - ID of newly created user	Cockpit Administrator
/user/ CockpitUserDelete	Deletes a cockpit user	deleteFromUAA userId username		Cockpit Administrator
/group/ GroupResourceAdd	Adds a resource to a group	groupId resourceId		Cockpit Resource Administrator
/group/ GroupResourceRemove	Removes a resource from a group	groupId resourceId		Cockpit Resource Administrator
/group/ GroupUserAdd	Adds a user to a group	groupId userId		Cockpit Resource Administrator
/group/ GroupUserRemove	Removes a user from a group	groupId userId		Cockpit Resource Administrator

5.1.7.2 Using Cockpit GET APIs

Details on SAP HANA cockpit APIs that provide information about registered resources and resource groups. The GET APIs don't create, delete, or change anything in the cockpit.

Each API with "Get" in its name (RegisteredResourcesGet, for example) is protected with both authentication and authorization checks; you must be authenticated as a SAP HANA cockpit user to invoke the APIs.

To call GET APIs programmatically, follow the instructions in [Managing Resources, Users, and Groups with the Cockpit APIs \[page 98\]](#) for getting a token and calling against the cockpit-landscape-svc. Look below for details on what the GET APIs do and which privileges you need to use them.

To call GET APIs from a browser, follow the instructions below.

When you call GET APIs from a browser, you cannot go directly to the cockpit-adminui-svc or cockpit-landscape-svc, since although the browser can send a GET it cannot support sending an authentication token. That's where the app-router comes in: it acts as a middleman, ensuring all browser requests that pass through it include an authentication token.

If you're using a browser to test a GET you must go through the app-router, using either the cockpit-admin-web-app or cockpit-web-app URL.

¹ See also the general responses described above under Success and Failure.

RegisteredResourcesGet

Returns information about the resources registered in SAP HANA cockpit.

Privileges required: COCKPIT_RESOURCE_ADMIN or COCKPIT_POWER_USER. The information returned depends on your role:

- If you have COCKPIT_RESOURCE_ADMIN, RegisteredResourcesGet returns information on all resources registered with this cockpit.
- If you have COCKPIT_POWER_USER, RegisteredResourcesGet returns information on all resources registered by you with this cockpit.

You can invoke RegisteredResourcesGet in two ways:

- Through the cockpit-admin-web-app port via the app-router. The call is redirected to the XSA sign-in page if it doesn't present an app-router cookie indicating authentication status. This form of invocation through the app-router is ideal for testing with a Web browser, but not ideal for programmatic calls.
- Against the cockpit-adminui-svc endpoint. Calls using this method must present a valid authentication token.

API Endpoints for RegisteredResourcesGet

cockpit-admin-web-app	/cp/admin/resource/RegisteredResourcesGet
cockpit-adminui-svc	/resource/RegisteredResourcesGet

RegisteredResourcesGet supports query parameters in OData format, for example, \$count, \$top=, \$skip=, \$orderby=.

When it succeeds, RegisteredResourcesGet returns HTTP status 200 and the result data. Otherwise, it returns an HTTP response code and text describing why the request failed – for example, 403, "Permission Denied."

A successful response is a JSON object in this format:

```
{
  "result": [
    {
      "BuildNumber": "1522210459",
      "CertificateHostName": "",
      "CollectionConfigurations": [],
      "Connections": [
        {
          "Host": "host.domain.com",
          "IsSAPControlAuthenticated": false,
          "PortType": "INSTANCE",
          "PortValue": 1,
          "Role": "MASTER",
          "SAPControlUserName": null
        }
      ],
      "CreatedBy": "COCKPIT_ADMIN",
      "DatabaseName": "DBNAME",
      "Designation": "CUSTOM",
      "EncryptedJDBC": false,
      "EncryptedSAPControl": false,
      "GroupCount": 1,
      "HardwarePlatform": "30A8S20Q0B",
      "Host": "host.domain.com",
      "HostName": "host.domain.com",
      "OSVersion": "SUSE Linux Enterprise Server 12.1",
    }
  ]
}
```

```

    "PatchLevel": 0,
    "Port": 1,
    "PortType": "INSTANCE",
    "ResIcon": "HANA_TENANTDB.gif",
    "ResKey": "RESKEY_HANA_MDB_TENANT",
    "ResValue": "HANA_MDB_TENANT",
    "ResourceDescription": "",
    "ResourceId": "ResourceId",
    "ResourceName": "DBNAME@SID",
    "ResourceOwnerDetail": "details",
    "ResourceOwnerEmail": "",
    "ResourceOwnerName": "",
    "ResourceUniqueId": "0f208532-8fcc-4aef-9bd1-40470aa03ad8",
    "SAPUser": null,
    "SSOEnabled": false,
    "SSOEnforced": false,
    "SSOSupported": true,
    "ServicePack": 30,
    "SystemName": "SID",
    "TechnicalUser": "SYSTEM",
    "ValidateServerCertificate": false,
    "Version": "2.00.030.00.1522210459 (hanaws, 2018.13.0)",
    "VersionMajor": 2,
    "VersionMinor": 0
  }
]
}

```

Each object in the array represents a registered resource. The object returned has these properties:

Property	Description
ResourceId	The internal ID the cockpit uses to identify this resource. Other cockpit APIs require this ID as a parameter.
ResourceUniqueId	The unique internal identifier of the resource. This value is either specified at registration time or read from the system being registered where possible. If not supplied by the caller or the resource a random value is generated but this is not guaranteed to be unique.
ResourceName	The name of the resource. For SAP HANA systems the default name takes the form <DB>@<SID> – for example, DB1@HA0.
ResourceDescription	
CreatedBy	The cockpit user who registered the resource. (This is not the same as the resource owner, below).The description of the resource optionally provided by the registering user at registration time.
SystemName	The system name. For SAP HANA systems this is the SID (HA0, for example).
DatabaseName	The database name, for example, DB1.
Designation	The designation or usage type of the resource. For SAP HANA systems these can be PRODUCTION, DEVELOPMENT, TESTING or CUSTOM.
Version	The resource's full version string.

Property	Description
VersionMajor	The major version of the resource as an integer (for example, 2)
VersionMinor	The minor version of the resource as an integer (for example, 0)
ServicePack	The service pack of the resource as an integer (for example, 20)
PatchLevel	The patch level of the resource as an integer.
BuildNumber	The build number of the resource.
ResKey	The description of the resource optionally provided by theA resource key for the type of the resource (for example "RESKEY_HANA_MDB_SYSTEM").
ResValue	The type of the resource (for example, HANA_SYSTEM).
ResIcon	A path to an icon used to render the resource type (not currently used)
OSVersion	The version of the operating system where the resource is running
HardwarePlatform	The platform where the resource is running (for example, VMware Virtual Platform).
ResourceOwnerName	Optional owner name. This is not necessarily the same person who registered the system and does not have to be a cockpit user. This is someone to contact with questions or problems.
ResourceOwnerEmail	Optional. The e-mail address of the owner.
ResourceOwnerDetail	Optional. Additional details on the resource owner. This might include work hours, location, or phone number.
SSOSupported	If this value is true, this resource supports single sign-on from the cockpit.
SSOEnabled	If this value is true, single sign-on from the cockpit is turned on.
SSOEnforced	If this value is true, single sign-on is the only available authentication method from the cockpit.
EncryptedJDBC	If this value is true, communication with this resource uses an encrypted database connection.
EncryptedSAPControl	If this value is true, communication with the host agent managing this resource uses TLS.
ValidateServerCertificate	If this value and EncryptedJDBC are both true, the client validates the server (resource) using pre-installed certificates.
CertificateHostName	If set the connection uses this value instead of the host name provided in the server certificate for encrypted database connections
TechnicalUser	The name of the technical user the cockpit uses for this resource.

Property	Description														
Host (& HostName)	The main host name of the resource. Duplicated in both fields for compatibility.														
PortType	The type of the main connection to the resource. Set to either "INSTANCE" or "SQL". If INSTANCE the port is an instance number. If SQL the port is a SQL port.														
Port	The instance number or SQL port (for example, the indexserver SQL port) of the main connection to the resource.														
GroupCount	The number of groups this resource belongs to.														
Connections	All possible connections to the resource (for scale-out systems or those with host aliases, for example). The main connection as dictated by Role is copied into the properties above. An array of connection objects in this form:														
	<table border="1"> <thead> <tr> <th>Property</th> <th>DescriptionDesc</th> </tr> </thead> <tbody> <tr> <td>Role</td> <td>The role of the connection: MASTER, MASTER_ALIAS, TENANT_MASTER, SLAVE, or STANDBY.</td> </tr> <tr> <td>PortType</td> <td>INSTANCE or SQL.</td> </tr> <tr> <td>PortValue</td> <td>Instance number or SQL port number.</td> </tr> <tr> <td>Host</td> <td>Host name for this connection.</td> </tr> <tr> <td>IsSAPControlAuthenticated</td> <td>If true, authenticated has been set (by the current user) for the SAP Control functionality.</td> </tr> <tr> <td>SAPControlUserName</td> <td>The name of the SAP Control user, if set.</td> </tr> </tbody> </table>	Property	DescriptionDesc	Role	The role of the connection: MASTER, MASTER_ALIAS, TENANT_MASTER, SLAVE, or STANDBY.	PortType	INSTANCE or SQL.	PortValue	Instance number or SQL port number.	Host	Host name for this connection.	IsSAPControlAuthenticated	If true, authenticated has been set (by the current user) for the SAP Control functionality.	SAPControlUserName	The name of the SAP Control user, if set.
Property	DescriptionDesc														
Role	The role of the connection: MASTER, MASTER_ALIAS, TENANT_MASTER, SLAVE, or STANDBY.														
PortType	INSTANCE or SQL.														
PortValue	Instance number or SQL port number.														
Host	Host name for this connection.														
IsSAPControlAuthenticated	If true, authenticated has been set (by the current user) for the SAP Control functionality.														
SAPControlUserName	The name of the SAP Control user, if set.														

GroupsForUserGet

Returns information about the resource groups (including the default groups Development, Production, and Test) that are visible to you.

You can invoke GroupsForUserGet in two ways:

- Through the cockpit-web-app port via the app-router. The call is redirected to the XSA sign-in page if it doesn't present an app-router cookie indicating authentication status. This form of invocation through the app-router is ideal for testing with a Web browser, but not ideal for programmatic calls.
- Against the cockpit-landscape-svc endpoint. Calls using this method must present a valid authentication token.

API Endpoints for GroupsForUserGet

cockpit-web-app	/cp/ls/group/GroupsForUserGet
cockpit-landscape-svc	/group/GroupsForUserGet

GroupsForUserGet supports query parameters in OData format, for example, \$count, \$top=, \$skip=, \$orderby=.

When it succeeds, GroupsForUserGet returns HTTP status 200 and the result data. Otherwise, it returns an HTTP response code and text describing why the request failed – for example, 403, "Permission Denied."

A successful response is a JSON object in this format:

```
{
  "d": {
    "results": [
      {
        "Description": "...",
        "Id": "123",
        "Name": "GROUPNAME",
        "__metadata": {
          "id": "https://host.domain.com:post/od/Groups(123L)",
          "type": "com.sap.hana.cockpit.persistence.odata.Group",
          "uri": "https://host.domain.com:post/od/Groups(123L)"
        }
      }
    ]
  }
}
```

Each object in the array represents a resource group. The object returned has these properties:

Property	Description
Type	An integer representing the type of the group: 1 = ALL group 2 = AUTO group (Production, Development, or Test) 3 = Group created by a user
Name	The name of the group.
Description	Optional. The description of the group provided by the user who created it.
CreatedBy	For Type 3 groups only. The cockpit user who created this resource group.
Id	For Type 3 groups only. The groupId of this group.
ResourceCount	The number of resources in this group.
RunWithAlert	The number of resources in this group that have active alerts.
NotRunning	The number of resources in this group that are not in the state RUNNING..

GroupResourcesGet

Returns information about the resources in a specified group that's visible to you. Only groups returned by GroupsForUserGet can be used as arguments for GroupResourcesGet.

GroupResourcesGet supports two forms of parameters: either a groupId obtained with GroupsForUserGet, or the groupDesignation of an automatic group. These are either PRODUCTION, DEVELOPMENT, TEST, CUSTOM or ALL. The ALL group represents every resource that you are authorized to see; it's a useful way for a non-admin user to get the list of all visible resources. If you specify a groupId you do not have access to, an error will occur.

You can invoke GroupResourcesGet in two ways:

- Through the cockpit-web-app port via the app-router. The call is redirected to the XSA sign-in page if it doesn't present an app-router cookie indicating authentication status. This form of invocation through the app-router is ideal for testing with a Web browser, but not ideal for programmatic calls.
- Against the cockpit-landscape-svc endpoint. Calls using this method must present a valid authentication token.

API Endpoints for GroupResourcesGet

cockpit-web-app	/cp/ls/group/GroupResourcesGet
cockpit-landscape-svc	/group/GroupResourcesGet

GroupResourcesGet requires one of these query parameters:

- ?groupId=<id> or
- ?groupDesignation=<string>

GroupResourcesGet supports query parameters in OData format, for example, \$count, \$top=, \$skip=, \$orderBy=.

When it succeeds, GroupsForUserGet returns HTTP status 200 and the result data. Otherwise, it returns an HTTP response code and text describing why the request failed – for example, 403, "Permission Denied."

A successful response is a JSON object in this format:

```
{
  "result": [
    {
      "AlertCountHigh": 2,
      "AlertCountMedium": 0,
      "Availability": 3,
      "AvailableGroups": {
        "_deferred": {
          "uri": "https://host.domain.com:port/pd/
ResourceOverviews(123L)"
        }
      },
      "BuildNumber": "1493036600",
      "Capacity": -1,
      "Connections": [
        {
          "Host": "host",
          "IsSAPControlAuthenticated": false,
          "PortType": "INSTANCE",
          "PortValue": 0,
          "Role": "MASTER",
          "SAPControlUserName": null
        }
      ],
      "DatabaseName": "",
      "Designation": "PRODUCTION",
      "GroupCount": 1,
      "Host": "host.domain.com",
      "HostName": "host.domain.com",
    }
  ]
}
```

```

    "IsAuthenticated": false,
    "IsAuthenticatedWithSSO": false,
    "PatchLevel": 9,
    "Performance": -1,
    "Port": 0,
    "PortType": "INSTANCE",
    "RemoteUserName": null,
    "ResValue": "HANA_SYSTEM",
    "ResourceDescription": "",
    "ResourceId": "123",
    "ResourceName": "ResourceName",
    "SAPControlAuthenticated": false,
    "SAPControlUser": "",
    "SSOEnabled": false,
    "SSOEnforced": false,
    "SSOSupported": null,
    "ServicePack": 122,
    "State": "UNKNOWN",
    "SystemName": "SystemName",
    "UserGroupCount": 1,
    "Version": "1.00.122.09.1493036600 (fa/hana1sp12)",
    "VersionMajor": 1,
    "VersionMinor": 0,
    "XSASupported": null
  }
]
}

```

Each object in the array represents a resource group. The object returned has these properties:

Property	Description
ResourceId	The internal ID the cockpit uses to identify this resource. Other cockpit APIs require this ID as a parameter.
ResourceName	The name of the resource. For SAP HANA systems the default name takes the form <DB>@<SID> – for example, DB1@HA0.
ResourceDescription	The description of the resource optionally provided by the registering user at registration time.
DatabaseName	The database name, for example, DB1.
SystemName	The system name. For SAP HANA systems this is the SID (HA0, for example).
Designation	The designation or usage type of the resource. For SAP HANA systems these can be PRODUCTION, DEVELOPMENT, TESTING or CUSTOM.
Version	The resource's full version string.
VersionMajor	The major version of the resource as an integer (for example, 2)
VersionMinor	The minor version of the resource as an integer (for example, 0)
ServicePack	The service pack of the resource as an integer (for example, 20)
PatchLevel	The patch level of the resource as an integer.

Property	Description
BuildNumber	The build number of the resource.
ResValue	The type of the resource (for example, HANA_SYSTEM).
SSOSupported	If this value is true, this resource supports single sign-on from the cockpit.
XSASupported	If this value is true, this resource has a running XS advanced server.
SSOEnabled	If this value is true, single sign-on from the cockpit is turned on.
SSOEnforced	If this value is true, single sign-on is the only available authentication method from the cockpit.
IsAuthenticated	If this value is true, the current user is authenticated with this resource.
IsAuthenticatedWithSSO	If this value is true, the current user is authenticated with this resource using single sign-on.
RemoteUserName	The database user currently used to authenticate with the resource.
Host (& HostName)	The main host name of the resource. Duplicated in both fields for compatibility.
State	The state of the resource (RUNNING or STOPPED, for example).
Availability	A score describing the availability of the resource.
Performance	A score describing the performance of the resource.
Capacity	A score describing the capacity of the resource.
AlertCountHigh	The number of high-level alerts currently active.
AlertCountMedium	The number of medium-level alerts currently active.
PortType	The type of the main connection to the resource. Set to either "INSTANCE" or "SQL". If INSTANCE the port is an instance number. If SQL the port is a SQL port.
Port	The instance number or SQL port (for example, the indexserver SQL port) of the main connection to the resource.
GroupCount	The number of groups this resource belongs to.
Connections	<p>All possible connections to the resource (for scale-out systems or those with host aliases, for example). The main connection as dictated by Role is copied into the properties above.</p> <p>An array of connection objects in this form:</p>
	Property
	DescriptionDesc

Property	Description
Role	The role of the connection: MASTER, MASTER_ALIAS, TENANT_MASTER, SLAVE, or STANDBY.
PortType	INSTANCE or SQL.
PortValue	Instance number or SQL port number.
Host	Host name for this connection.
IsSAPControlAuthenticated	If this value is true, authenticated has been set (by the current user) for the SAP Control functionality.
SAPControlUserName	The name of the SAP Control user, if set.

5.1.8 Using XS CLI Commands to Troubleshoot the Cockpit

If you encounter issues with SAP HANA cockpit, you can use `xs cli` commands to view services logs and application status.

You can execute the `xs cli` commands on the machine where the cockpit is installed, using the `<sid>adm` account, or remotely using the XSA Client. For complete details on logging into the SAP HANA XS advanced runtime console and on the XS CLI: Application Management commands, see *SAP HANA Developer Guide for SAP HANA XS Advanced Model*.

Viewing Logs of Various Services

You can investigate potential issues by viewing the log file of a specific service with the command `xs logs <APP>`, where `<APP>` is the name of the application whose log-file details you want to display.

If you encounter issues with...	View the application log for...
Resource registration, resource group management, cockpit user management, or other cockpit manager issues	<ul style="list-style-type: none"> cockpit-admin-web-app cockpit-admin-ui-svc cockpit-persistence-svc
Displaying or retrieving data, or issues related to specific SAP HANA cockpit applications	<ul style="list-style-type: none"> cockpit-web-app cockpit-hdbui-svc cockpit-hdb-svc cockpit-landscape-svc cockpit-persistence-svc

If you encounter issues with...

Collections (for example, if you notice that information in the Aggregate Health Monitor or in the alerts does not seem accurate)

View the application log for...

- cockpit-collection-svc
- cockpit-hdb-svc
- cockpit-persistence-svc

→ Tip

In the log of cockpit-collection-svc, if you see: `A collection could not be submitted for execution because the worker thread pool is exhausted`, then consider increasing the collection worker thread pool in the data collection settings through the cockpit manager. Increase the threads incrementally, rechecking the log each time, until the issue is resolved.

Viewing Application Status

After viewing log files, you can also look at application status with the `xs apps` command. Ensure that the following services are in the STARTED state, and that instances are up and running:

- hrtt-service
- sqlanz-svc
- sqlanz-ui hrtt-core
- sapui5_fesv2
- cockpit-persistence-svc
- cockpit-hdb-svc
- cockpit-collection-svc
- cockpit-hdbui-svc
- cockpit-landscape-svc
- cockpit-web-app
- cockpit-adminui-svc
- cockpit-admin-web-app

5.2 SAP HANA Studio

The SAP HANA studio runs on the Eclipse platform and is both a development environment and administration tool for SAP HANA.

Administrators can use the SAP HANA studio, for example, to start and stop services, to monitor the system, to configure system settings, and to manage users and authorizations. The SAP HANA studio accesses the servers of the SAP HANA database by SQL. Developers can use the SAP HANA studio to create content such as modeled views and stored procedures. These development artifacts are stored in the repository, which is part of the SAP HANA database. The SAP HANA studio is developed in Java and based on the Eclipse platform.

The SAP HANA studio presents its various tools in the form of perspectives. Database administration and monitoring features are available primarily within the SAP HANA Administration Console perspective.

Additional perspectives include the SAP HANA Modeler perspective and the SAP HANA Development perspective. For more information about these perspectives, see the *SAP HANA Developer Guide (For SAP HANA Studio)* and the *SAP HANA Modeling Guide (For SAP HANA Studio)*.

i Note

Depending on how you installed the studio, all features may not be available. During installation, you can specify which features you require depending on your role. For system administration, only the feature SAP HANA Studio Administration is necessary. For more information, see *SAP HANA Studio Features* in the *SAP HANA Studio Installation and Update Guide*.

Updating the SAP HANA Studio

To ensure that you are working with the most recent version of the SAP HANA studio, you need to check regularly for updates. You can update the SAP HANA studio using several methods. For example, you can use SAP HANA Software Lifecycle Manager, or you can set up a central update site.

For more information, see the *SAP HANA Studio Installation and Update Guide*.

5.2.1 Open the SAP HANA Administration Console

To access the database administration and monitoring features of the SAP HANA studio, you open the SAP HANA Administration Console perspective.

Procedure

1. From your file explorer, start `hdbstudio.exe`.
2. On the *Welcome* page, choose *Open SAP HANA Administration Console*.

Results

The SAP HANA Administration Console opens. The *Systems* view is open by default. This view is the central access point for performing system-specific administration and monitoring activities. From this view, you can access the other views and editors used for administration.

i Note

Once you have closed the *Welcome* page, you can always change from another perspective to the SAP HANA Administration Console perspective by choosing **Window > Open Perspective > SAP HANA Administration Console** or by choosing the  *SAP HANA Administration Console* button in the perspective switcher in the upper-right corner of the screen.

Related Information

Screen Areas of the SAP HANA Administration Console [page 115]

Editors and Views of the SAP HANA Administration Console [page 116]

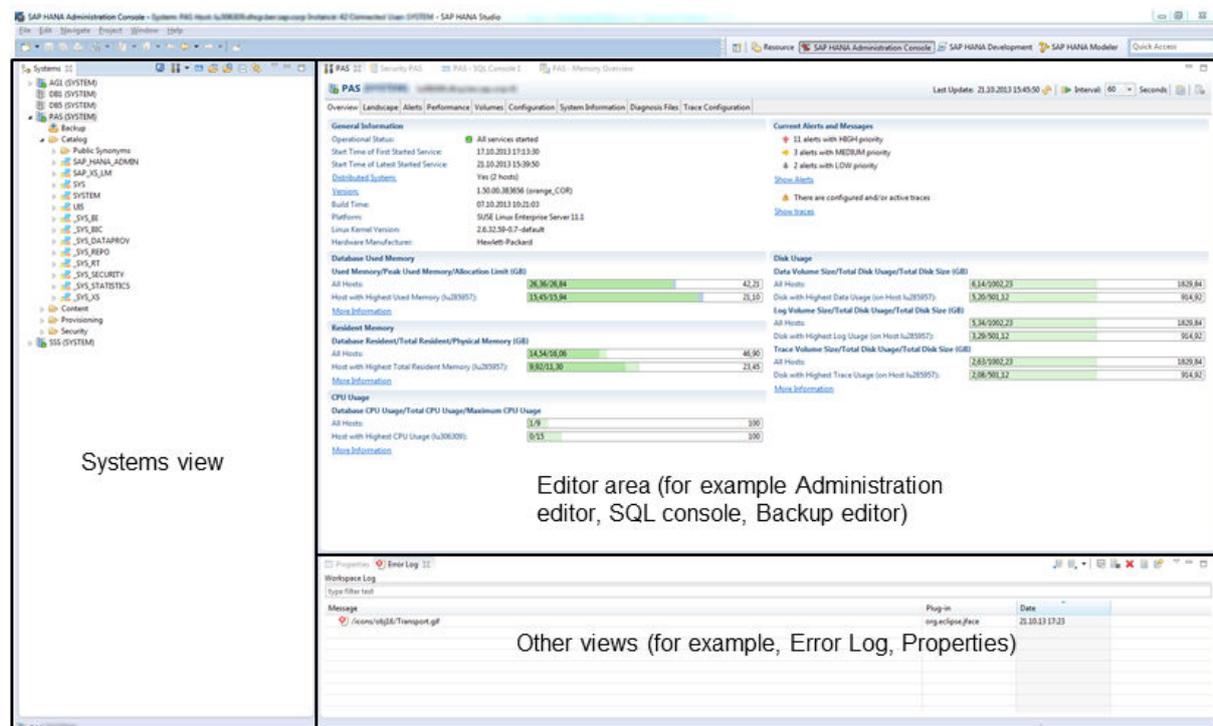
Systems View [page 121]

5.2.1.1 Screen Areas of the SAP HANA Administration Console

The database administration and monitoring features of the SAP HANA studio are presented in the SAP HANA Administration Console perspective according to a default screen layout.

The following figure shows the screen areas of the SAP HANA studio with the SAP HANA Administration Console perspective open:

Screen Areas of the SAP HANA Administration Console



The following is a brief overview of the various screen areas:

Screen Area	Description
Main menu and main toolbar	<p>The main menu contains standard Eclipse functions.</p> <p>The main toolbar is located beneath the main menu. The contents of this toolbar change based on the active perspective. Items in the toolbar might be enabled or disabled based on the state of either the active view or editor. The perspective switcher is an additional toolbar normally located on the top-right of the screen, next to the main toolbar. It allows quick access to perspectives that are currently open. It also has a button that can open new perspectives.</p>
Editor area	<p>Each perspective has editors for editing and browsing resources. Editors are displayed as tabs in the editor area. Several editors can be open at the same time.</p> <p>Important editors available in the SAP HANA Administration Console include:</p> <ul style="list-style-type: none"> • The System Monitor • The Administration editor • The Backup editor • The Security editor • The SQL console
Views	<p>Views support editors and provide alternative presentations as well as ways to navigate the information in the SAP HANA studio. Important views available in the SAP HANA Administration Console include:</p> <ul style="list-style-type: none"> • Systems, which is the central access point for performing system administration and monitoring activities • Error Log, which contains error and information messages • Properties, which shows the detailed properties of the active resource (for example, the SAP HANA system selected in the Systems view) <p>To open a view, from the main menu, choose ▶ Window ▶ Show View ▾.</p>

For more information about the Eclipse platform, see the Eclipse documentation.

5.2.1.2 Editors and Views of the SAP HANA Administration Console

Several editors and views are available in the Administration Console for the administration and monitoring of SAP HANA databases.

The following table describes the main system-level editors and views available in the [Administration Console](#) and how to access them. Other editors are available for specific resources (for example users, roles, tables and so on).

View/Editor	Description	How to Open
Systems	The Systems view provides you with a hierarchical view of all the SAP HANA systems managed in the SAP HANA studio and their contents. It is the central access point for performing system-specific administration and monitoring activities using the other available editors.	The Systems view is open by default when you open the Administration Console. If it is closed, you can open it from the main menu by choosing Window > Show View > Systems .
System Monitor	The System Monitor is an editor that provides you with an overview of all your SAP HANA systems at a glance. From the System Monitor , you can drill down into the details of an individual system in the Administration editor.	In the toolbar of the Systems view, choose the  button.
Administration	The Administration editor the main tool for performing administration and monitoring activities.	You can access the Administration editor in several ways: <ul style="list-style-type: none"> From the Systems view toolbar, choose the  Open Default Administration button. In the Systems view, double-click the system. In the context menu of the Systems view, choose Configuration and Monitoring > Open Administration.
Administration Diagnosis Mode	The Administration editor diagnosis mode allows you to monitor and perform emergency operations on systems to which either no SQL connection is available or the SQL connection is overloaded.	The Administration editor opens automatically in diagnosis mode in the following situations: <ul style="list-style-type: none"> When you open the Administration editor for a system that cannot be reached by SQL When you initiate the start, stop, or restart of a system <p>You can also open the Administration editor in diagnosis mode from the Systems view toolbar by choosing the  Open Diagnosis Mode button.</p>
Backup	The Backup editor is the main tool for performing administration and monitoring activities related to backup.	You can access the Backup editor in several ways: <ul style="list-style-type: none"> Expand the system in the Systems view and choose the  Backup entry In the context menu of the Systems view, choose Backup and Recovery > Open Backup Console.

View/Editor	Description	How to Open
Security	<p>The Security editor is the main tool for managing the following aspects of security administration:</p> <ul style="list-style-type: none"> • Password policy • Auditing • Data volume encryption 	<p>You can access the Security editor in several ways:</p> <ul style="list-style-type: none"> • Expand the system in the Systems view and choose the  Security entry • In the context menu of the Systems view, choose  Security  Open Security Console .
SQL Console	<p>Some tasks may require you to work with SQL statements, for example, certain administration tasks can only be performed using SQL. You can enter, execute, and analyze SQL statements in the SQL console.</p>	<p>You can access the SQL console in several ways:</p> <ul style="list-style-type: none"> • From the Systems view toolbar, choose the  button. • In the context menu of the Systems view, choose Open SQL Console.

Related Information

[Systems View \[page 121\]](#)

[System Monitor \[page 177\]](#)

[Administration Editor \[page 359\]](#)

[Reference: Backup Console \(SAP HANA Studio\) \[page 1396\]](#)

[Execute SQL Statements in SAP HANA Studio \[page 118\]](#)

5.2.2 Execute SQL Statements in SAP HANA Studio

You can execute SQL statements in the SAP HANA studio using the SQL console.

Prerequisites

- You have added the system in the [Systems](#) view. For more information, see [Add an SAP HANA System](#).
- You have the required privileges to perform the operation. For more information about privileges, see the [User Authorization](#).
- (Optional) You have customized the behavior of SQL statement execution in the SQL console. You can do this in the SAP HANA studio preferences ( [SAP HANA](#)  [Runtime](#)  [SQL](#) ). For more information, see [SAP HANA Studio Administration Preferences](#).

Procedure

1. Open the SQL console:
 - a. Select the system in the *Systems* view.
 - b. From the toolbar, choose the  (*Open SQL console for selected system*) button.

The SQL console displays the connected system and user in the header. If you opened the SQL console from a specific catalog object, the schema is also displayed.

To connect to a different system from within the SQL console, choose the  (*Choose Connection*) button in the toolbar in the top-right of the editor and choose another system.

2. Enter the SQL statement or statements.

The following rules apply:

- You can write SQL syntax elements in either upper or lower case.
- You can add any number of spaces and line breaks.
- To force the system to distinguish between upper/lower-case letters in database object names (such as table names), enter the name between double quotation marks: "My_Table"
- To comment out a line, use - - (double hyphens) at the start of the line.
- To use name completion, press the key combination `CTRL` + `SPACE`.
This opens a list from which you can choose schema and table names, SQL keywords, and user-defined templates.

Note

You define templates in the preferences (▶ *SAP HANA* ▶ *Runtime* ▶ *Templates* ▶).

- Enter multiple SQL statements, separated by the configured separator character, semicolon (;) by default.
3. Execute the statement(s) in one of the following ways:
 - In the context menu, choose *Execute*.
 - Choose the  *Execute* button in the toolbar.
 - Press `F8`.

If you have entered several statements, you can execute them individually by simply highlighting the statement and executing. If you do not highlight an individual statement, all statements are executed.

By default, statements are prepared before execution. You can disable statement preparation in the preferences.

Results

The *Result* tab appears with the statement's results. Multiple *Result* editors may open depending on the statement(s) executed.

Note

To prevent performance issues with the SAP HANA studio, by default only 50 *Result* editors can open. Once this number is reached, statement execution stops.

Information about statement execution is displayed in the lower part of the screen, for example:

```
Started: 2013-11-27 14:22:16
Statement 'SELECT * FROM "PUBLIC"."M_CS_TABLES"'
Successfully executed in 260 ms 932 µs (server processing time: 258 ms 868 µs)
Fetched 583 row(s) in 16 ms 602 µs (server processing time: 11 ms 278 µs)
```

i Note

SAP HANA implements a prefetch mechanism that returns the head of the result together with the execute command. By default, 32 rows are prefetched. This means that if the result set is smaller than the number of rows prefetched, the subsequent fetch command can simply take the rows from the prefetch buffer in the client library without any further processing on the server.

Related Information

[Add an SAP HANA System \[page 122\]](#)

[Systems View \[page 121\]](#)

[SAP HANA Studio Administration Preferences \[page 138\]](#)

[User Authorization \[page 738\]](#)

5.2.3 Managing SAP HANA Systems in SAP HANA Studio

Before you can start working with SAP HANA systems in the SAP HANA studio, you must first add and connect to them. Additional features allow you to manage systems efficiently and conveniently in the studio.

Related Information

[Systems View \[page 121\]](#)

[Add an SAP HANA System \[page 122\]](#)

[Add an SAP HANA System with a Different User \[page 126\]](#)

[Link a Centrally-Stored Archive of SAP HANA Systems \[page 128\]](#)

[Log Off From/Log On To an SAP HANA System \[page 129\]](#)

[Disable Password Storage in Eclipse Secure Store \[page 130\]](#)

[Organize the Systems View Using Folders \[page 131\]](#)

[Search for SAP HANA Systems \[page 132\]](#)

[Export a List of SAP HANA Systems \[page 132\]](#)

[Import a List of SAP HANA Systems \[page 133\]](#)

[Disable Default Filtering of Schemas \[page 134\]](#)

[Start the SAP HANA Studio with Immediate System Logon \[page 135\]](#)

[Configure TLS/SSL for SAP HANA Studio Connections \[page 136\]](#)

5.2.3.1 Systems View

The *Systems* view provides you with a hierarchical view of all the SAP HANA systems managed in the SAP HANA studio and their contents. It is the central access point for performing system-specific administration and monitoring activities.

You can use the SAP HANA studio to manage both SAP HANA systems with and without tenant databases. In the *Systems* view, systems that do not support tenant databases are identified solely by their SID. Tenant databases are identified by their database name and SID (<database_name>@<SID>). The system database can also be registered in the *Systems* view; it is identified by database name and SID (SYSTEMDB@<SID>). For more information about adding systems and tenant databases, see *Add an SAP HANA System*.

The *Systems* view comprises the following elements:

- A hierarchical view of all your systems/databases and their contents. For administration and monitoring purposes, the following are the most important elements:
 - The *Catalog* folder, which contains all activated database objects, grouped by schema
 - The *Security* folder, which contains all database users and activated roles

Double-clicking the top-level system entry in the hierarchical view opens it in the Administration editor.

From the hierarchical view, you can also access the Backup (📁) and Security (🔑) editors.

- A toolbar that provides you with quick access to several editors and functions.
- A context menu that provides you quick access to a range of system-specific functions.

Related Information

[Toolbar Options in the Systems View \[page 121\]](#)

[Add an SAP HANA System \[page 122\]](#)

5.2.3.1.1 Toolbar Options in the Systems View

The *Systems* view toolbar provides you with quick access to several editors and functions.

Icon	Option	Description
	<i>Add System...</i>	Opens the <i>Add System</i> dialog in which you can create and configure a connection to a system
	<i>Add System Archive Link...</i>	Opens the <i>Add System Archive Link</i> dialog in which you can add a link to a centrally-stored archive of SAP HANA systems
	<i>Open System Monitor</i>	Opens the System Monitor to see an overview of all systems in the Systems view
	<i>Open Default Administration</i>	Opens the Administration editor for the selected system
	<i>Open Diagnosis Mode</i>	Opens the Administration editor for the selected system in diagnosis mode

Icon	Option	Description
	<i>Open SQL Console</i>	Opens the SQL console for the selected system
	<i>Find System</i>	Opens the <i>Find System</i> dialog in which you can search for a system in the Systems view
	<i>Collapse All</i>	Collapses the tree expansion state of all systems in the Systems view
	<i>Link with Editor</i>	Toggles whether the entry selected in the <i>Systems</i> view is linked to the active editor. When this option is selected, changing the active editor will automatically update the selected system.
	<i>View Menu</i>	Provides menu items that allow you to sort or filter the contents of the <i>Systems</i> view

Related Information

[Add an SAP HANA System \[page 122\]](#)

[Link a Centrally-Stored Archive of SAP HANA Systems \[page 128\]](#)

[System Monitor \[page 177\]](#)

[Administration Editor \[page 359\]](#)

[Troubleshooting an Inaccessible or Unresponsive SAP HANA System \[page 684\]](#)

[Execute SQL Statements in SAP HANA Studio \[page 118\]](#)

5.2.3.2 Add an SAP HANA System

To work with and manage an SAP HANA system using the SAP HANA studio, you must create and configure a connection to it.

Prerequisites

- The relevant ports in your firewall are open.
- If you want to secure communication between the SAP HANA server and the SAP HANA studio using the Secure Sockets Layer (SSL) protocol, you have configured the server for SSL and imported the trust store file that contains the server root certificate into either the Java keystore or your user keystore on the client.
- You must have a database user, or the necessary infrastructure for Kerberos-based user authentication must be in place.

Procedure

1. From the *Systems* view toolbar, choose  (*Add System...*).
The *System* wizard opens.
2. Enter the required system information:
 - System connection properties (required)
 - Host name
 - Instance
 - Mode

For more information about what to enter for these properties, see *System Connection Properties*.
 - Description (optional)
This is the description of the system that you want to appear next to the system name in the *Systems* view.
 - Folder (optional)
If you are organizing your systems in the *Systems* view using folders and have already created folders, choose the folder to which you want to add the system.
 - Locale (optional)
This setting specifies the language of objects created in the SAP HANA repository.
3. Choose *Next*.
4. Choose the authentication type for user logon to the database:
 - If you are integrating the SAP HANA studio into a Kerberos-based single sign-on environment, choose *Authentication by current operating system user*.
 - If you are implementing user name/password authentication, choose *Authentication by database user* and enter the database user name and password. You can choose to have your password stored in the Eclipse secure storage so that you do not have re-enter it every time you open the studio.
5. Indicate whether you want to use a secure connection to the system by choosing *Connect using SSL*.

i Note

You must select this option to be able to modify the SSL connection properties (step 9).

6. Configure the connection of the SAP start service (`sapstartsrv`) to the system.
An HTTP connection to the system using the SAP start service is automatically enabled. You can choose to disable this connection if it is not required.

i Note

If you disable this connection, administrative actions that require operating system access are not possible in the SAP HANA studio (for example, stopping and starting the system, performing a recovery, or opening the Administration editor in diagnosis mode).

If you want the SAP start service to communicate with the system via a secure connection, choose *Use HTTPS*.

7. Choose *Next*.
8. Optional: Modify the following advanced connection properties for your system:

Option	Description
Option	JDBC connection parameter(s)
Auto-Reconnect	Auto-reconnect option If you select this option, the SAP HANA studio automatically reconnects if the connection to the system fails.

9. Optional: Configure SSL communication:

- To have the identity of the server validated during connection, choose [Validate SSL Certificate](#). The server's public-key certificate is validated against the root certificate stored in the trust store. If you want to override the system host name specified in the server certificate, enter a host name with a defined certificate.
- [Use user key store as trust store](#)
The Java SSL property `trustStore` specifies the trust store containing the certificate used to validate that the server certificate is issued by a trusted entity. Each user can import certificates into a user keystore in Java using the `keytool` command line tool (part of the JRE installation). The user keystore is located in the home directory of the current operating system user. The file name is `.keystore`. The set of root certificates delivered with the JRE from well-known issuers (for example, Verisign, Thawte, Deutsche Telekom) is used when this option is not selected.

10. Choose [Finish](#).

The system is added in the [Systems](#) view. The system entry displays the following information:

- SID
In the case of a multitenant database container, the database name is indicated before the SID, either `SYSTEMDB@<SID>` or `<database>@<SID>`
- Connected user
- System usage if it is a production system
- System or database description if available

Information about system availability and user connection status are indicated by icons.

Results

You can now access the system or database in the SAP HANA studio.

Related Information

[User Authentication and Single-Sign On \[page 715\]](#)

[Provisioning Users \[page 775\]](#)

[System Connection Properties \[page 125\]](#)

[Configure System Usage Type \[page 304\]](#)

[Configure TLS/SSL for SAP HANA Studio Connections \[page 136\]](#)

[Creating and Configuring Tenant Databases \[page 189\]](#)

5.2.3.2.1 System Connection Properties

To connect to an SAP HANA system or database in the SAP HANA studio, you must specify its connection properties (host, instance, and mode).

Property	Description
Host name	<p>Fully qualified domain name (FQDN) of the host on which the system is installed</p> <ul style="list-style-type: none">Multi-host system If you are adding a multi-host system, specify the master host. You do not have to enter all host names explicitly as they are determined automatically. If the master host becomes unavailable, the connection is automatically established through one of the other hosts. Hosts that are added to the system later are also detected automatically. <div data-bbox="667 790 1394 987"><p>→ Tip</p><p>Once you have finished adding the system, you can see all available hosts in the system properties. Right-click the system in the <i>Systems</i> view and choose <i>Properties</i>. All hosts are listed on the <i>Hosts Used to Connect</i> tab of the <i>Database User Logon</i> page.</p></div> <div data-bbox="667 1003 1394 1173"><p>i Note</p><p>The host name of the server that hosts the database must be accessible from the client on which the SAP HANA studio is running, even if you add the system using its IP address.</p></div> <ul style="list-style-type: none">Tenant databases If you are adding a tenant database in a tenant database system, specify the FQDN of the system host.
Instance	Instance number of the system
Mode	<p>Whether the system you are connecting to is a single-container system or multiple-container system</p> <p>If you are connecting to a system with tenant databases, then you must further specify the specific database that you want to connect to:</p> <ul style="list-style-type: none">The system databaseA tenant database

5.2.3.3 Add an SAP HANA System with a Different User

If you want to work with an SAP HANA system using several database users, you can create a connection to the system or database in the SAP HANA studio with the credentials of different users.

Prerequisites

- The system has already been added once in the *Systems* view.
- If you want to secure communication between the SAP HANA server and the SAP HANA studio using the Secure Sockets Layer (SSL) protocol, you have configured the server for SSL and imported the trust store file that contains the server root certificate into either the Java keystore or your user keystore on the client.
- The relevant ports in your firewall are open.
- You must have a database user, or the necessary infrastructure for Kerberos-based user authentication must be in place.

Procedure

1. In the *Systems* view, right-click the system and choose *Add System with Different User*.
2. Choose the authentication type for user logon to the system:
 - If you are integrating the SAP HANA studio into a Kerberos-based single sign-on environment, choose *Authentication by current operating system user*.
 - If you are implementing user name/password authentication, choose *Authentication by database user* and enter the database user name and password. You can choose to have your password stored in the Eclipse secure storage so that you do not have re-enter it every time you open the studio.
3. Indicate whether you want to use a secure connection to the system by choosing *Connect using SSL*.

i Note

You must select this option to be able to modify the SSL connection properties (step 7).

4. Configure the connection of the SAP start service `sapstartsrv` to the system.
An HTTP connection to the system using the SAP start service is automatically enabled. You can choose to disable this connection if it is not required.

i Note

If you disable this connection, administrative actions that require operating system access are not possible in the SAP HANA studio (for example, stopping and starting the system, performing a recovery, or opening the Administration editor in diagnosis mode).

If you want the SAP start service to communicate with the system via a secure connection, choose *Use HTTPS*.

5. Choose *Next*.

6. Optional: Modify the following advanced connection properties for your system:

Option	Description
Option	JDBC connection parameter(s)
Auto-Reconnect	Auto-reconnect option If you select this option, the SAP HANA studio automatically reconnects if the connection to the system fails.

7. Optional: Configure SSL communication:

- To have the identity of the server validated during connection, choose [Validate SSL Certificate](#). The server's public-key certificate is validated against the root certificate stored in the trust store. If you want to override the system host name specified in the server certificate, enter a host name with a defined certificate.
- [Use user key store as trust store](#)
The Java SSL property `trustStore` specifies the trust store containing the certificate used to validate that the server certificate is issued by a trusted entity. Each user can import certificates into a user keystore in Java using the `keytool` command line tool (part of the JRE installation). The user keystore is located in the home directory of the current operating system user. The file name is `.keystore`. The set of root certificates delivered with the JRE from well-known issuers (for example, Verisign, Thawte, Deutsche Telekom) is used when this option is not selected.

8. Choose [Finish](#).

The system is added in the [Systems](#) view. The system entry displays the following information:

- SID
In the case of a tenant database system, the database name is indicated before the SID, either `SYSTEMDB@<SID>` or `<database>@<SID>`
- Connected user
- System usage if it is a production system
- System description if available

Information about system availability and user connection status are indicated by icons.

Results

You can now access the system or database in the SAP HANA studio.

Related Information

[User Authentication and Single-Sign On \[page 715\]](#)

[Provisioning Users \[page 775\]](#)

[Configure System Usage Type \[page 304\]](#)

[Configure TLS/SSL for SAP HANA Studio Connections \[page 136\]](#)

5.2.3.4 Link a Centrally-Stored Archive of SAP HANA Systems

To allow users who work in the SAP HANA studio to connect to multiple SAP HANA systems, you can manage a list of all systems in a centrally-accessible archive. Users can then simply link to this archive.

Prerequisites

An XML file containing a list of all SAP HANA systems and their connection information exists at a centrally-accessible location, for example, a network file server.

You can create this file by exporting a list of systems from your installation of the SAP HANA studio to the required location.

Context

A centrally-stored archive of SAP HANA systems allows you to deploy system information to all users of the SAP HANA studio, for example, developers, content modelers, and other administrators. It avoids users having to obtain the connection details of all systems individually and then having to add them all individually. In addition, if you change the central file, for example to add new systems or change the host of an existing system, you can ensure that users always have up-to-date system access.

Procedure

1. From the *Systems* view toolbar, choose  (*Add System Archive Link...*).
2. Specify the following link details:
 - Link name
 - Path to the system archive containing the system information
 - Optional: A folder in the *Systems* view
3. Choose *Finish*.

Results

The system archive appears in the *Systems* view as a link node (). By expanding the link node, you can see all the systems contained within.

To be able to access a system in the system archive, the password of the connecting user specified in the system properties must be available in the user's local Eclipse secure storage. If this is not the case, you must log on to the system.

i Note

The system archive file does **not** contain user passwords.

As the system archive is only linked, note the following:

- Systems are not added to the user's local workspace.
- Users cannot edit the connection properties of systems in the system archive.
- Users cannot change the order or hierarchical structure of systems in the system archive.

Related Information

[Export a List of SAP HANA Systems \[page 132\]](#)

5.2.3.5 Log Off From/Log On To an SAP HANA System

In the SAP HANA studio, you can log off from an SAP HANA system and close all connections to the system. To be able to connect to system again, you must log on.

Procedure

- To log off from a system right-click it in the *Systems* view and choose *Log Off*. All open connections to the system are closed, and in the *Systems* view, the system appears disabled. No information regarding its operational status is available; you cannot expand it and browse its contents.

i Note

Editors connected to the system at the time of log-off may close as a result. If an editor contains any unsaved work, you will be prompted to save it first.

- To log on to a system, simply double click it in the *Systems* view or from the context menu, choose *Log On*. If your password is saved in the Eclipse secure store, you are logged on to the system immediately and can connect to it again. If have disabled the storing of passwords in the Eclipse secure store, you must re-enter your password.

Related Information

[Disable Password Storage in Eclipse Secure Store \[page 130\]](#)

5.2.3.5.1 User Logon Behavior on SAP HANA Studio Startup

Whether or not you are logged on to your SAP HANA systems when you start the SAP HANA studio depends on whether or not you were logged on when you closed the studio.

If you logged off from a system before closing the studio, you are still logged off and must log on explicitly. If you were logged on when you closed the studio, you are logged on automatically. This is the default behavior.

However, you can change this behavior so that no automatic logon takes place when the studio is started: explicit logon is always required.

To do so, choose **► Preferences ► SAP HANA ► Global Settings ►** and deselect the option *Restore logged-on/ logged-off status of systems on startup*.

i Note

Automatic logon on studio startup can only take place if the connecting user's password is stored in the Eclipse secure store. If it is not, explicit logon is always required.

5.2.3.6 Disable Password Storage in Eclipse Secure Store

When an SAP HANA system is added in the SAP HANA studio, the user can choose to store his or her password in the Eclipse secure storage. To improve security, you can disable this password storage. Users must then log on to the system every time they open the studio.

Prerequisites

You are logged on to the computer on which the SAP HANA studio is installed as either the root user (Linux) or local administrator (Windows).

Context

The Eclipse secure storage stores user passwords securely to disk on the SAP HANA studio client. To connect to a system in the SAP HANA studio, the user does not have to enter his or her password; the stored password is used. This behavior may not be desired for security reasons in some cases, for example:

- To prevent individuals from being able to access systems using another user's credentials
This is possible if several users share the computer on which the SAP HANA studio is installed.
- To prevent users from locking their accounts
This is possible if a user's password for a system has expired but the old password is stored in the secure store. The user may lock their account due to too many failed logon attempts.

Procedure

Disable password storage by specifying the command `-noPwdStore` in one of the following ways:

- As a start-up parameter of `hdbstudio.exe` (for example, in the program shortcut properties of a Windows installation)
- As a parameter in the `hdbstudio.ini` configuration file

Results

User passwords cannot be stored in the Eclipse secure storage. When the SAP HANA studio is opened, systems appear in a logged-off state in the *Systems* view.

To connect to the system, the user must log on to it by choosing *Log On* from the context menu and then entering his or her password. The password is stored temporarily for the duration of the session only. The session ends when the user closes either the SAP HANA studio or the individual system by choosing *Log Off* from the context menu.

5.2.3.7 Organize the Systems View Using Folders

If you add several SAP HANA systems in the *Systems* view, you can define a folder structure to organize them.

Procedure

1. From the main menu, choose **► New > Folder ◄**.
2. Enter a folder name.
3. In the *Systems* view, move your system to the new folder using drag and drop.
4. Repeat this procedure until you have added all your systems.

Results

Once folders have been created, you can assign any new systems to a folder when you add them.

5.2.3.8 Search for SAP HANA Systems

If you have a large number of systems registered in the *Systems* view, you can search for a specific system to access it more quickly.

Procedure

1. From the *Systems* view toolbar, choose the  (*Find System*) button.
2. Enter a search string.
You can also use * or ? as wildcards.
Matching systems are displayed.
3. Select the system you were searching for.
You can select several systems in the search results by pressing the **CTRL** key while selecting. You can use this, for example, to mark duplicate systems.
4. Choose whether you want to open the selected system in the Administration editor and/or the SQL console.

Results

The system opens in the Administration editor and/or SQL console. If you did not select either of these options, the system is only highlighted in the *Systems* view.

5.2.3.9 Export a List of SAP HANA Systems

You can export a list of your SAP HANA systems from the SAP HANA studio as an XML file and then import it into another instance of the SAP HANA studio or use it as system archive to which other users can link.

Procedure

1. From the main menu, choose **File** > *Export...*
2. Expand the *SAP HANA* folder and choose *Landscape*.
3. Choose *Next*.
4. Select the systems you want to export and enter a target file location.
5. Choose *Finish*.

Results

The list of systems and their properties (name, description, host name, instance, and so on) is exported as an XML file to the specified location.

Related Information

[Link a Centrally-Stored Archive of SAP HANA Systems \[page 128\]](#)

5.2.3.10 Import a List of SAP HANA Systems

You can import a list of SAP HANA systems that you previously exported from another instance of the SAP HANA studio.

Procedure

1. From the main menu, choose **File > Import...**
2. Expand the *SAP HANA* folder and then choose *Landscape*.
3. Choose *Next*.
4. Choose *Browse...* and select the file containing the list of systems that you want to import.
5. Select the folder into which you want to import the file.
6. Choose *Finish*.

Results

The systems are added in the *Systems* view of the SAP HANA studio.

To be able to access the systems, the password of the connecting user specified in the system properties must be available in the user's local Eclipse secure storage. If this is not the case, you must log on to the system.

i Note

The file containing the list of systems does **not** contain user passwords.

5.2.3.11 Disable Default Filtering of Schemas

Users with the system privilege DATA ADMIN and/or CATALOG READ, for example database administrators, may not see all schemas in the [Systems](#) view of the SAP HANA studio since a default filter is applied.

Context

In the [Systems](#) view of the SAP HANA studio, users only see those schemas for which at least one of the following criterion applies:

- The user has at least one object privilege on the schema.
- The user has at least one object privilege on at least one object in the schema.
- The user owns at least one object in the schema.

i Note

For all privilege checks, not only privileges directly granted to the user but also privileges granted to one of his or her roles (or to roles in these roles) are considered.

As a result, users with the system privilege DATA ADMIN and/or CATALOG READ cannot see all available schemas.

If, as a database administrator, you need to see all available schemas, you must disable this default schema filter.

Procedure

1. In the [Systems](#) view, right-click [Catalog](#) and choose [Filters...](#)
The [Filter for Schema](#) dialog box opens.
2. Select [Display all schemas](#).
3. Optional: Specify a filter pattern to reduce the number of schemas displayed.
This is useful if the total number of schemas exceeds the number of displayable items in the tree (configured under [► Preferences > Catalog > ►](#)). If this is the case, then you will not see all schemas at once and will have to browse.
4. Save and apply the filter by choosing [OK](#).

Results

Schemas are displayed filtered according the specified filter pattern.

5.2.3.12 Start the SAP HANA Studio with Immediate System Logon

The SAP HANA studio program accepts command line parameters that allow you to specify the system to be connected to immediately on startup. This can be useful to system administrators, as well as to other programs that call the SAP HANA studio.

Prerequisites

You have a database user in the SAP HANA system that you want to log on to.

Procedure

1. Launch the SAP HANA studio from its installation directory passing the following start parameters:

Option	Description
-h	Host name
-n	Instance number
-u	User name

Note

User names containing special characters that represent conjunction or redirection characters in the command line program must be enclosed in double quotation marks ("..."), regardless of where the special character appears in the user name.

Example

Windows:

- `hdbstudio.exe -h hana1 -n 02 -u DBADMIN`
- `hdbstudio.exe -h hana1 -n 02 -u "&test"`

Linux

- `hdbstudio -h hana1 -n 02 -u DBADMIN`
- `hdbstudio -h hana1 -n 02 -u "&test"`

Mac OS:

- `open -a /Applications/sap/hdbstudio.app --args -h hana1 -n 02 -u DBADMIN`
- `open -a /Applications/sap/hdbstudio.app --args -h hana1 -n 02 -u "&test"`

The SAP HANA studio opens.

2. If prompted, enter your user password.

Results

The system is added in the [Systems](#) view (if it is not already there), and you are logged on.

5.2.3.13 Configure TLS/SSL for SAP HANA Studio Connections

Secure communication between the SAP HANA studio and the SAP HANA database using the Transport Security Layer (TLS)/Secure Sockets Layer (SSL) protocol.

Prerequisites

- You have configured the SAP HANA database for secure client-server communication over JDBC/ODBC. For more information, see *SSL Configuration on the SAP HANA Server* in the *SAP HANA Security Guide*.
- You have added the SAP HANA system in the SAP HANA studio.

Context

The SAP HANA studio communicates with the SAP HANA database via the JDBC client interface. The client-side configuration of the SAP HANA studio uses Java TLS/SSL properties.

Procedure

1. Using the keytool command line tool, import the truststore file that contains the server root certificate into either the Java keystore or your personal user keystore.

By default, the SAP HANA studio client validates server certificate(s) against the root certificate stored in the Java keystore of the running VM (virtual machine). This keystore is part of the Java installation and is located in the Java home directory under `${JAVA_HOME}/lib/security/cacerts` (Linux) or `%JAVA_HOME%/lib/security/cacerts` (Windows).

However, it is not recommended that you store the root certificate in this keystore, but in your personal user keystore instead. The user keystore is located in the home directory of the current operating system user. The file name is `.keystore`.

2. Enable and configure TLS/SSL secure communication between the SAP HANA studio and the server:

In the SAP HANA studio, open the system's properties and choose [Connect Using SSL](#).

This corresponds to setting the Java SSL property `encrypt` to `true`.

3. Configure how the identity of the server is to be validated during connection (server-side authentication):

- a. In the system's properties dialog, choose the *Additional Properties* tab.
- b. If you want server certificate(s) to be validated using the default truststore, choose *Validate SSL Certificate*.

This corresponds to setting the Java SSL property `validateCertificate` to **true**.

When an TLS/SSL connection is established, the host name in the certificate being connected to and the host name in the server certificate must match. This may not always be the case. For example, in a single-host system, if a connection is established from the SAP HANA studio on the same host as the SAP HANA server, a mismatch would arise between the host named in the certificate (fully qualified host name) and the host used to establish the connection (`localhost`)*.

You can override the host name specified in the server certificate by entering a host name with a defined certificate in the *Override Host Name Certificate* field. This corresponds to setting the Java SSL property `hostNameInCertificate`.

- c. If you want the server certificate to be validated using the user's keystore and not the default Java keystore, choose *Use user keystore as trust store*.

This corresponds to changing the value of the Java SSL property `trustStore`.

Note

If you do not have a working public key infrastructure (PKI), you can also suppress server certificate validation entirely by selecting neither of these options (*Validate SSL Certificate* or *Use user keystore as trust store*). However, this is not recommended.

4. Optional: If the identity of the client is to be validated by the SAP HANA server (client certificate validation), perform the following additional steps:
 - a. In the *Additional Properties* tab of the system properties, specify the path to the user keystore that contains your private key, as well as the pass phrase required to access this file.
 - b. Enable validation of the client's identity on the server by changing the parameter `[communication] sslValidateCertificate` in the `global.ini` file to **true**.

You can do this on the *Configuration* tab of the Administration editor.

- c. Import the client root certificate into the server truststore used for client-server communication.

If you manage client certificates directly in the database (recommended), this means importing the certificate into the certificate store and adding it to the certificate collection with the purpose *SSL*.

Results

In the *Systems* view, a lock icon appears next to the system name () , indicating that SSL communication is active.

Related Information

[Add an SAP HANA System \[page 122\]](#)

[Managing Client Certificates \[page 900\]](#)

5.2.4 SAP HANA Studio Administration Preferences

The preferences of the SAP HANA studio include many options for customizing the features of the SAP HANA Administration Console.

To open the preferences of the SAP HANA studio, choose **Window > Preferences**. The preferences related to SAP HANA perspectives are all available under *SAP HANA*.

The following preferences pages contain administration-related settings:

- [Administration \[page 138\]](#)
- [Global Settings \[page 139\]](#)
- [Runtime \[page 139\]](#)
- [Table Viewer \[page 142\]](#)

Administration

Administration

Option	Description
Show user-defined SQL statements on the <i>System Information</i> tab	If you select this option, user-defined SQL statements contained in the specified XML file are displayed on the <i>System Information</i> tab of the Administration editor. You can also change the default location and name of the XML file.

Backup Editor

Option	Description
Number of SQL objects to retrieve	This setting determines the number of backups displayed on the <i>Backup Catalog</i> tab of the Backup editor.
Refresh interval in seconds	This setting determines the refresh interval of the <i>Overview</i> tab of the Backup editor.
Connection timeout in seconds	This setting sets a timeout for the connection to the backup editor. If the specified timeout is exceeded, no further attempt is made to establish a connection to the Backup editor. As a consequence, the Backup editor is not displayed in the <i>Systems</i> view. The other information about the SAP HANA database is still shown in SAP HANA studio.

Global Settings

Global Settings

Option	Description
Restore logged-on/logged-off status of systems on startup	<p>This option determines whether or not you are automatically logged on to systems registered in the Systems view. By default, if you were logged on when you closed the studio (and your password is saved in the Eclipse secure store), you are logged on automatically on restart. Similarly, if you logged off before closing the studio, you are not logged on restart and you must actively log on.</p> <p>If you deselect this option, you must always log on after restart.</p>
Request confirmation before a user is deleted	<p>When a user is deleted, all dependent objects are also deleted. Select this option if you want a confirmation message to appear before a user is deleted.</p>
Show Management Console for <code>hdbcons</code>	<p>If you select this option, an additional tab Console is available in the Administration editor. You can execute <code>hdbcons</code> commands directly in this console.</p> <div style="border: 1px solid #ccc; padding: 5px;"><p>⚠ Caution</p><p>Technical expertise is required to use <code>hdbcons</code>. To avoid incorrect usage, use <code>hdbcons</code> only with the guidance of SAP HANA development support.</p></div>

Runtime

Catalog

Option	Description
Fetch all database catalog objects	<p>By default the SAP HANA studio fetches a limited number of catalog objects when folders in the Systems view such as Tables and Views are opened.</p> <p>If you select this option, all catalog objects are loaded in the corresponding folder. This may affect system performance as it may take some time to fetch all database catalog objects.</p>
Number of database catalog objects to display	<p>If you do not select the Fetch all database catalog objects option, you can specify the maximum number of catalog objects to be fetched. If the number of available objects exceeds the number specified here, the message Object limit reached appears.</p> <p>The default number is 1,000.</p>
Show table comment before table name (Modeler)	<p>If you select this option, a table's description appears before its name in the Systems view if the SAP HANA Modeler perspective is active.</p>

Common

Option	Description
Confirm saving editors	<p>If you select this option, the system displays a confirmation dialog when an editor is closed with content that was not saved.</p>

Option	Description
Autosaving of SQL Console Content	If you select this option, the content of SQL console sessions is saved automatically when the SAP HANA studio is closed. No dialog requesting the user to save is displayed.
<ul style="list-style-type: none"> Save content when SAP HANA Studio is closed Content save interval ... minutes 	Additionally, it is possible to have the content saved at a specified interval. If the SAP HANA studio is closed unexpectedly, the last version can be recovered.
Copy options:	These are formatting options for copying content from the table editor.
<ul style="list-style-type: none"> Data separator Tab separated Align copied values with space Copy cell in editor by using <code>[CTRL] C</code> Copy editor content with column header 	
Representation of null value	This option specifies the character used to display NULL values
Database identifier upper case	This option specifies that the IDs of database objects can be entered only in uppercase letters.
Default action for database tables:	This setting specifies which view of a table is opened when it is double-clicked in the <i>Systems</i> view: its definition or its content.
<ul style="list-style-type: none"> Show content Show definition 	
Table Distribution Editor	This setting specifies the maximum number of tables that are displayed when you show table distribution.
Maximum Number of Tables Displayed	
Result	
Option	Description
Limit for LOB columns (bytes)	This option specifies the maximum number of bytes that are loaded from the database for one large object (LOB) column.
Limit for zoom (bytes)	This option specifies the maximum number of bytes that the SAP HANA studio displays when you zoom the LOB column in the result table in the <i>Result</i> tab of the SQL console.
Append exported data to file	If you select this option, then when you export the result table to a file, the system attaches the content of the current result table to the existing file content.
Display character byte value as hex	If you select this option, data of the data type CHAR BYTE is displayed as hexadecimal digits. If you do not select this option, this data is displayed in binary format.
Format values	If you select this option, country-specific formatting is applied (for example, numeric values or dates).
Display duration result row fetch	If you select this option, you can see in the SQL console how long it took to fetch one row of a result set.
Maximum number of rows displayed in result	This option specifies the maximum number of rows fetched from the database and displayed in the result table of the <i>Result</i> tab.

Option	Description
Enable zoom of LOB columns	<p>You must select this option if you want to be able to zoom LOB columns in the result table of the Result tab. You can zoom an LOB column by right-clicking and choosing .</p> <p>Note that if you zoom an LOB column, it is automatically closed after 15 minutes or when the <i>Result</i> tab is closed.</p>

Note

The options available under *Result* relate to the display of results following execution of a SELECT statement the SQL console.

SQL

Option	Description
Stop batch SQL statement execution if an error occurs	If you select this option, then when you execute a series of SQL statements separated by comment characters, the system stops the execution when an error occurs.
Clear SQL console log before SQL statement execution	If you select this option, the log from the last SQL statement is deleted before the next SQL statement is executed.
Close results before SQL statement execution	If you select this option, then when you execute an SQL statement in the SQL console, all old results tabs in the same SQL console session are closed.
Display time of statement execution start	If you select this option, you can see in the SQL console the time at which statement was executed.
Display duration of failed statements	If you select this option, you can see in the SQL console how long a statement took to execute in the SAP HANA studio even if the statement failed.
Connection parameters for SQL console: <ul style="list-style-type: none"> • Auto-commit mode • Isolation level • Confirm change of connection 	<ul style="list-style-type: none"> • Auto-commit mode: <ul style="list-style-type: none"> ○ If on, the system performs all COMMIT actions automatically ○ If off, you have to enter COMMIT statements explicitly. • The isolation level determines how the system implicitly controls locking and versioning of database objects. • Confirm change of connection In the SQL console, you can change the SQL connection you are working on. When you change a connection, cursors may be closed or transactions may be rolled back. If you select this option, a change of SQL connection must first be confirmed.

SQL console settings

Command separator	This option specifies the separator for SQL statements in the SQL console.
Maximum number of characters for multiple statement execution	When you enter multiple statements in the SQL console for execution, the content must be parsed and the individual statements for execution recognized. However, if there is too much content, out-of-memory situations or a long parse time may result. When the number of characters specified with this option is reached, parsing does not take place and the content is executed as a single statement.
Number of tables for table name completion	This option specifies the number of tables that are displayed in the list when you use name completion in the SQL console.

Option	Description
Number of open <i>Result</i> editors	The results of statement execution may be returned in multiple <i>Result</i> editors. Once the number of open <i>Result</i> editors specified with this option is reached (by default 50), statement execution stops. We recommend that you do not increase the default value of this option as it may cause performance issues in the studio.

Templates

Option	Description
Name	Word to be completed when you press the key combination <code>CTRL</code> + <code>SPACE</code> . You can create more than one template with the same name. If more than one template exists for one word, the system displays a list.
Context	Editor in which you can use the template.
Description	Template description
Auto insert	If on, the code assist automatically inserts the template if it is the only proposal available at the cursor position.

i Note

The options available under *Templates* always refer to the editor that is currently open.

Table Viewer

Table Viewer

Option	Description
Show gridlines	Use these options to customize the appearance of list displays in the Administration editor, for example, the list of files on the <i>Diagnosis Files</i> tab.
Alternating colored rows	

5.3 SAP HANA Hardware Configuration Check Tool for Tailored Data Center Integration

The SAP HANA HW Configuration Check Tool allows you to check the interoperability of SAP HANA with your existing enterprise storage, network and server in production environments.

In addition to SAP HANA as standardized and highly optimized appliance, SAP offers the opportunity to run the SAP HANA server with a customer's preferred storage and network solutions. This option enables you to reduce hardware and operational costs through the reuse of existing hardware components and operational processes. Here an SAP HANA server means the exact same bill of material as the certified SAP HANA appliance but without storage. Certified SAP HANA servers that can be used in a TDI deployment are listed in the PAM and the SAP HANA Hardware Directory.

The SAP HANA HW Configuration Check Tool is a framework that provides tests and reports for new single host and scale out systems to determine if the hardware you intend to use meets the minimum performance criteria required to run SAP HANA in production use.

Caution

The test should only be used before going into production. It should only be used on production machines when this has first been requested by SAP support.

Related Information

[SAP Note 1943937](#)

[Product Availability Matrix for SAP HANA](#)

[Certified and Supported SAP HANA Hardware Directory](#)

5.3.1 Install the SAP HANA Hardware Configuration Check Tool

The SAP HANA Hardware Configuration Check Tool allows you to measure the performance of your hardware components to ensure they meet the criteria for running SAP HANA.

Prerequisites

- Check which version of the tool you require. See SAP Note 1943937.
- Check SAP Note 2235581 for supported operating system versions. Check the SAP Notes listed in SAP Note 2235581 for OS-related settings.
- You should be able to run this tool set as the root user.
- When using the SAP HANA HW configuration check tool set to evaluate a distributed landscape the binaries should be available from a shared directory so every server can execute it. This means that the same version of the tool set must be installed on each server in the system.
- Check that you have libnuma1 v2 installed.
- Ensure that you have already exchanged SSH keys between the different hosts so that the workflow can take place without passwords.
- Check that hostname resolution works in both directions for fully qualified domain names.
- SAPCAR is needed to extract the binaries.

Context

Follow the instructions in SAP Note 1943937 to download the latest version of the tool as a SAR file from the SAP Software Download Center.

We recommend that you put the binaries in a shared location, for example in a directory parallel to your main SAP HANA installation directory like `/hana/shared/`. This avoids any potential problems with sharing the binaries for distributed tests. Make sure that the same version is installed on every server.

Procedure

1. Copy the SAR file HWCCT.SAR to the system hosting your SAP HANA server.
2. Install the tool by executing this command:

```
SAPCAR -xf HWCCT.SAR hwcct
```

Results

A new directory `/hana/shared/hwcct/` is created.

Detailed information on configuring tests and using the tool is contained in SAP Note 1943937.

Related Information

[SAP Software Download Center](#) 

[SAP Note 1943937](#) 

[SAP Note 2235581](#) 

5.4 SAP Solution Manager for SAP HANA Administration

SAP Solution Manager allows you to manage your business applications throughout their entire lifecycle. You can integrate SAP HANA into an overall operations concept supported through SAP Solution Manager, as of release 7.1, SPO5.

SAP HANA is often used in conjunction with other SAP business applications. For example, an SAP ERP system might call accelerators on SAP HANA to speed up business processes, or a product such as SAP Business Warehouse is deployed on the SAP HANA database. If you are using SAP HANA in such a context, then you must manage your business application in addition to administering the in-memory database. This is best done using an integrated approach.

SAP provides you with the SAP Solution Manager application management platform as part of your maintenance agreement. You can use it to manage your business applications throughout their entire lifecycle. As of release 7.1, SP05, SAP Solution Manager supports integration with SAP HANA. You can optimize your operational processes using this combined approach. One example is root cause analysis. Let's assume you have detected a problem in an application that is deployed on SAP HANA or calls an SAP HANA accelerator. In this case, you first have to find out whether the problem is caused by the application or by the SAP HANA database. SAP Solution Manager allows you to trace a process across all included components (from the user interface to the database) to locate the source of the problem. Then, detailed analysis speeds up your resolution process.

Other examples of how SAP HANA and SAP Solution Manager can be valuably integrated in the area of system operation are the processes for monitoring and change control. If your business application is still running on a traditional database, even integrated database administration might be relevant.

Related Information

[Controlling Change \(example of integration with Solution Manager\) \[page 160\]](#)

[Analyzing the Root Cause \(example of integration with Solution Manager\) \[page 160\]](#)

[Central Monitoring and Administration with SAP Solution Manager \[page 159\]](#)

[Connecting SAP Solution Manager to SAP HANA \[page 145\]](#)

[SAP Help Portal: SAP Solution Manager Documentation](#)

[SAP Support Portal: SAP Solution Manager for SAP HANA \(various resources including videos\)](#)

[SAP Support Portal: SAP Solution Manager Early Knowledge Transfer](#)

[SAP Community: SAP Solution Manager \(wikis, blogs, news\)](#)

[SAP Community \(Technical Operations wiki\): SAP HANA Managed System Setup in SAP Solution Manager](#)

[SAP Community \(Technical Operations\): Troubleshooting Guide for SAP Solution Manager](#)

[Solution Manager documentation for Landscape Management Database \(LMDB\)](#)

5.4.1 Connecting SAP Solution Manager to SAP HANA

Configure a connection to SAP HANA in SAP Solution Manager.

If you want to use capabilities of SAP Solution Manager, you have to make sure that the two systems know each other. Prerequisite for this is the registration of the SAP HANA system in the System Landscape Directory (SLD). From there, SAP Solution Manager gets the information that the SAP HANA system exists. The communication between the systems is based on a central agent infrastructure. The pre-configured agents are delivered by SAP and deployed on the SAP HANA appliance by the hardware partner.

The configuration of the connection itself is done as part of the basic configuration of SAP Solution Manager. In the guided procedure for Managed Systems Configuration you just need to set up the correct connection, assign the right agents, enter some parameters, create required users, and do a few more configurations. After this, you can start the collaboration of SAP HANA and SAP Solution Manager.

Some of the processes in SAP Solution Manager require additional configuration to specify how they should handle the SAP HANA database. For example, you have to specify in system monitoring which metrics you

want to control, or you have to define your transport landscape (development system to quality assurance system to production system) for change control.

Related Information

[Configuring, Working with and Administering System Landscape Directory \(SAP NetWeaver\)](#)

[SAP Community: System Landscape Directory Overview](#)

[Solution Manager documentation for Landscape Management Database \(LMDB\)](#)

[SAP HANA Operations with SAP Solution Manager](#)

[SAP HANA Managed System Setup in SAP Solution Manager](#)

[SAP Note 1747682](#)

5.4.1.1 Configuring an SAP HANA System to Connect to the System Landscape Directory (SLD)

You can use the SAP HANA database lifecycle manager to configure the connection parameters for the central System Landscape Directory (SLD) system.

The System Landscape Directory (SLD) serves as a central information repository for your system landscape. Data suppliers collect and send system data to SLD on a regular basis. The SLD data supplier for SAP HANA systems is implemented within the name server of the SAP HANA system. However, to enable the data collection process for your SAP HANA system, you must first configure the system's connection to the SLD. Note that the SAP HANA database lifecycle manager provides only the functionality to configure the connection to the SLD, the actual registration is performed automatically by the SLD data supplier afterward.

Related Information

[Configure SLD Registration Using the Graphical User Interface \[page 147\]](#)

[Configure SLD Registration Using the Command-Line Interface \[page 148\]](#)

[Configure SLD Registration Using the Web User Interface \[page 150\]](#)

[Using the SAP HANA Platform LCM Tools \[page 921\]](#)

5.4.1.1.1 Configure SLD Registration Using the Graphical User Interface

You can configure an SAP HANA system to connect to the System Landscape Directory (SLD) using the SAP HANA database lifecycle manager (HDBLCM) resident program in the graphical user interface.

Prerequisites

- The SAP HANA system has been installed or updated with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on with the required root user or system administrator user `<sid>adm` credentials.

Context

When an SAP HANA system is connected to the SLD, it can report its status and provide details and information about the system itself. For more information, see SAP Note 1673424 and SAP Note 1649323.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdbclm
```

By default, `<sapmnt>` is `/hana/shared`.

2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

```
./hdbclmgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

3. Select *Configure System Landscape Directory Registration* from the activity options. Then select *Next*.
4. Define the required parameters. Then select *Next*.

Field Name	Description
<i>SLD Host Name</i>	Specifies the name of the host where the SLD system is installed.
<i>SLD Port</i>	Specifies the standard HTTP access port of the SLD.
<i>SLD User Name</i>	Specifies the user of the SLD system. It must be a user that already exists on the host where the SLD system is running.
<i>SLD Password</i>	Specifies the password for the SLD system.

Field Name	Description
<i>Use HTTPS</i>	Specifies whether or not to use HTTPS.

- Review the summary, and select *Run* to finalize the configuration.

Next Steps

After you have configured the connection parameters, you can manually push the registration of the SAP HANA system in the central SLD system instead of waiting for the SAP HANA system to be registered automatically at a later point in time.

To do so, as a `<sid>adm` user, execute the following command:

```
/usr/sap/hostctrl/exe/saposcol -b | sldreg -connectfile /usr/sap/<SID>/SYS/global/slddest.cfg -stdin -oldtransferdtd
```

Related Information

[Using the SAP HANA Platform LCM Tools \[page 921\]](#)

5.4.1.1.2 Configure SLD Registration Using the Command-Line Interface

You can configure an SAP HANA system to connect to the System Landscape Directory (SLD) using the SAP HANA database lifecycle manager (HDBLCM) resident program in the command-line interface.

Prerequisites

- The SAP HANA system has been installed or updated with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on with the required root user or system administrator user `<sid>adm` credentials.

Context

When an SAP HANA system is connected to the SLD, it can report its status and provide details and information about the system itself. For more information, see SAP Note 1673424 and SAP Note 1649323.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdbclm
```

By default, <sapmnt> is /hana/shared.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

- o `./hdbclm`

and enter the index of the `configure_sld` action, or

- o Start the tool with the `configure_sld` action specified:

```
./hdbclm --action=configure_sld
```

3. Define the required parameters.

Field Name	Description
<i>SLD Host Name</i>	Specifies the name of the host where the SLD system is installed.
<i>SLD Port</i>	Specifies the standard HTTP access port of the SLD.
<i>SLD User Name</i>	Specifies the user of the SLD system. It must be a user that already exists on the host where the SLD system is running.
<i>SLD Password</i>	Specifies the password for the SLD system.
<i>Use HTTPS</i>	Specifies whether or not to use HTTPS.

For more information about parameters for SLD configuration, see Related Information.

4. Review the summary, and select `y` to finalize the configuration.

Results

You have configured the SAP HANA system to connect to the System Landscape Directory (SLD). The registration itself is performed by the SLD data supplier service.

This configuration task can also be performed in batch mode and using a configuration file. For more information about the available configuration methods, see *Using the SAP HANA Platform LCM Tools*.

i Note

When using the command line, the options can be set interactively during configuration only if they are marked as interactive in the help description. All other options have to be specified in the command line. To call the help, in the `hdbclm` directory of the SAP HANA system, execute the following command:

```
./hdbclm --action=configure_sld --help
```

❖ Example

The following example configures the registration in an SLD system interactively. The connection from the SLD Data Supplier service to the SLD is over HTTPS.

```
./hdblcm --action=configure_sld --sld_hostname=mysap --sld_port=50000 --  
sld_username=SLDuser --sld_password=SLDpassword -https
```

Next Steps

After you have configured the connection parameters, you can manually push the registration of the SAP HANA system in the central SLD system instead of waiting for the SAP HANA system to be registered automatically at a later point in time.

To do so, as a `<sid>adm` user, execute the following command:

```
/usr/sap/hostctrl/exe/saposcol -b | sldreg -connectfile /usr/sap/<SID>/SYS/  
global/slddest.cfg -stdin -oldtransferdtd
```

5.4.1.1.3 Configure SLD Registration Using the Web User Interface

A connection to the System Landscape Directory (SLD) can be configured using the SAP HANA database lifecycle manager Web user interface.

Prerequisites

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
 - Microsoft Windows
 - Internet Explorer - Version 9 or higher
If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing [Tools > Compatibility View Settings](#) .
 - Microsoft Edge
 - Mozilla Firefox - Latest version and Extended Support Release
 - Google Chrome - Latest version

- SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
- Mac OS - Safari 5.1 or higher

i Note

For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

Context

When an SAP HANA system is connected to the SLD, it can report its status and provide details and information about the system itself. For more information, see SAP Note 1673424 and SAP Note 1649323.

Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: <code>https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html</code>

i Note

The URL is case sensitive. Make sure you enter upper and lower case letters correctly.

SAP HANA cockpit	<ol style="list-style-type: none"> 1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. <code>https://<host_FQDN>:<port></code>
-------------------------	---

i Note

FQDN = fully qualified domain name

2. Drill down on the name of the system from *My Resources* or from a group.
3. The links in *Platform Lifecycle Management* each launch additional functionality, giving you expanded capabilities for managing the resource.

2. Select the *Configure System Landscape Directory Registration* tile.
3. Specify values for the following fields:

Field Name	Description
<i>SLD Host Name</i>	Specifies the name of the host where the SLD system is installed.
<i>SLD Port</i>	Specifies the standard HTTP access port of the SLD.

Field Name	Description
<i>SLD User Name</i>	Specifies the user of the SLD system. It must be a user that already exists on the host where the SLD system is running.
<i>SLD Password</i>	Specifies the password for the SLD system.
<i>Use HTTPS Connection</i>	Specifies whether or not to use HTTPS.

- Review the summary, and select *Run* to finalize the configuration.

Next Steps

After you have configured the connection parameters, you can manually push the registration of the SAP HANA system in the central SLD system instead of waiting for the SAP HANA system to be registered automatically at a later point in time.

To do so, as a `<sid>adm` user, execute the following command:

```
/usr/sap/hostctrl/exe/saposcol -b | sldreg -connectfile /usr/sap/<SID>/SYS/global/slddest.cfg -stdin -oldtransferdtd
```

Related Information

[SAPUI5 Developer Guide](#)

[SAP Note 1673424](#)

[SAP Note 1649323](#)

[Add an SAP HANA System \[page 122\]](#)

[Using the SAP HANA Platform LCM Tools \[page 921\]](#)

5.4.1.2 Change the Default SLD Data Supplier Configuration

The System Landscape Directory (SLD) is the central directory of system landscape information relevant for the management of your software lifecycle. Data suppliers collect and send system data to SLD on a regular basis. The SLD data supplier for SAP HANA systems is implemented within the name server.

Prerequisites

- The SLD is configured.
For more information, see SAP Note 1018839 and the introductory section *Configuring an SAP HANA System to Connect to the System Landscape Directory*.

- You have the system privilege INIFILE ADMIN.

Context

For SAP HANA systems, the name server contains the SLD data supplier. It is configured by default to automatically transfer data to the SLD on a regular basis using the `sldreg` executable. Data is transferred in XML format in a file named `sldreg.xml`. You can change the default settings if required by modifying the `nameserver.ini` configuration file; for example, it may not be necessary to send data to the SLD frequently if your landscape is stable, or you may need to change the default save locations of the configuration and log files.

You can edit configuration files in the SAP HANA cockpit or SAP HANA studio.

i Note

Links to additional resources related to SAP Solution Manager are listed under Related Links.

Procedure

1. Navigate to the `nameserver.ini`.
2. Add a new section `sld` and add those parameters whose default value you want to change.

The following table lists the possible parameters and their default values.

i Note

Under normal circumstances, you will not need to change the default values. It should only be necessary, for example, for testing purposes or if requested as part of a support inquiry.

Key	Meaning	Default Value	Note
<code>enable</code>	Activates or deactivates the SLD data supplier	<code>true</code>	Allowed values are <code>true</code> , <code>false</code> . Re-enabling this parameter triggers a new generation of <code>sldreg.xml</code> and sending to the SLD system.

Key	Meaning	Default Value	Note
enable_virtddbhome	In a system replication landscape, enables the SAP_IdenticalDatabaseSystem association in SLD between the virtual database and the physical database configured with the parameters sldvirtddbhome and sldsystemhome in the system_landscape_hostname_virtualization section of the global.ini file	Undefined	<p>Allowed values are true and false.</p> <p>This parameter facilitates the registration of an SAP HANA system replication landscape in SAP Solution Manager.</p> <p>With the sr_register command, this parameter is set to true in the primary system and to false in the secondary system(s).</p> <p>In the event of a takeover (sr_takeover), the value of this parameter is set to true in the new primary system and to false in the original primary system, thus changing the association of virtual databases to the physical databases in the new primary.</p> <p>For more information, see <i>Configuring SAP HANA for System Replication Technical Scenario in SAP Solution Manager</i>.</p>
Interval	Specifies the frequency (in seconds) with which the sldreg.xml file is generated. If a newly-generated document is the same as the previous one, it is not sent to the SLD.	300	<p>It does not make sense to enter small positive values or negative values.</p> <p>If you enter 0 or a negative value, data is transferred to the SLD only once.</p> <p>Enter a value without a "1000 separator" (for example, 1899, not 1,899 or 1.899), otherwise it is interpreted as 0.</p>
force_interval	Specifies how often (in seconds) the sldreg.xml file must be sent to the SLD, even if the file has not changed.	43200	
configpath	Specifies the location of the folder that contains the configuration file sldest.cfg This file is a parameter for the call to sldreg.	/usr/sap/ <sid>/SYS/ global	Example: /usr/sap/MPW/SYS/global

Key	Meaning	Default Value	Note
xmlpath	Specifies where the file <code>sldreg.xml</code> is generated and where the <code>sldreg.log</code> log file is written <code>sldreg.log</code> is the log file of <code>sldreg</code> , and both files are parameters for the call to <code>sldreg</code> .	<code>/usr/sap/</code> <code><sid>/</code> <code>HDB<id>/</code> <code><currenthost</code> <code>>/trace</code>	Example: <code>/usr/sap/LRH/HDB42/velber1cm1/trace</code>
lmstructurepath	Specifies where the <code>lm_structure</code> directory is created	<code>hana/shared/</code> <code><sid>/</code> <code>lm_structure</code>	

Results

The transfer of data will take place in line with your new settings.

Note

If errors occur in the transfer of data to the SLD, you can check the log file `sldreg.log` and the database trace for the name server with trace components `SLDCollect`, `SLDConfig`, and `SLDSend`.

Related Information

[Configuring SAP HANA System Properties \(INI Files\) \[page 291\]](#)

[Database Trace \(Basic, User-Specific, and End-to-End\) \[page 667\]](#)

[Configuring SAP HANA for System Replication Technical Scenario in SAP Solution Manager \[page 155\]](#)

[SAP Note 1018839](#)

[SAP Note 2082466](#)

5.4.1.3 Configuring SAP HANA for System Replication Technical Scenario in SAP Solution Manager

To model an SAP HANA system replication landscape in SAP Solution Manager as a technical scenario, you must configure SAP HANA to send the correct landscape data to the Landscape Management Database (LMDB) via the System Landscape Directory (SLD). SAP Solution Manager uses the system information managed in the LMDB.

System replication is a mechanism for ensuring the high availability of SAP HANA systems, as well as disaster recovery. Through the continuous replication of data from a primary to a secondary system, system replication facilitates rapid failover in the event of a disaster. Production operations can be resumed with minimal

downtime. With multitier system replication, a third system is attached to the first secondary making it a replication chain of three systems. For more information about system replication, see the section on availability and scalability in this guide.

Such a system replication landscape can be modeled as a technical scenario in SAP Solution Manager. The SAP HANA SLD supplier sends the required system landscape information to the LMDB based on the following landscape configuration settings in SAP HANA.

i Note

The following configuration must be in place before you enable system replication on the primary system and register secondary systems.

Landscape Configuration Settings

Define Virtual Databases (sldvirtddbhome)

A virtual database is the logical host name under which an SAP HANA database can always be reached for the purposes of **takeover**. This ensures that application systems remain associated with the correct database in SLD when a secondary system takes over from the primary system.

In each SAP HANA system in the system replication landscape, configure a virtual database at the system level, or if desired for each individual tenant database, by setting the parameter

[system_landscape_hostname_virtualization] sldvirtddbhome in the `global.ini` file of the system database. This will result in each virtual database being registered in SLD as an `SAP_HDBSystem` object (`<database_name>.DBTypeForSAP.HDB.SystemHome.<sldvirtddbhome>`).

Note the following:

- If you don't configure a value for each tenant database in the `global.ini`, the value configured at the system level applies. If you don't configure a value at the system level, the value of the `SLDVIRTDDBHOME` parameter in the default system profile (`hana/shared/<sid>/profile/DEFAULT.PFL`) is used; or if this value hasn't been set, the host name from the landscape properties file (`hana/shared/<sid>/lm_structure/landscapeVariables.properties`).
- Configure a virtual database for either all tenants or none. If some tenants are configured and some are not, the fallback values mentioned above are not used for those tenants without a configured value. In this case, the data supplier does not deliver this tenant completely and an error is written to the error log.
- Changing the `sldvirtddbhome` parameter in the `global.ini` file does not require a restart, whereas changing the default profile does.

Define Physical Databases (sldsystemhome)

A physical database is the logical host name used to identify an SAP HANA database independently of internal host names for the purposes of **auto-host failover**. Since the host names of master and standby hosts are different, configuring a single physical database ensures that the same database is registered in SLD after a failover as before.

In each SAP HANA system in the system replication landscape, you configure a unique physical database at the system level, or if desired for each individual tenant database, by setting the parameter

[system_landscape_hostname_virtualization] sldsystemhome in the `global.ini` file of the

system database. This will result in each physical database being registered in SLD as an `SAP_HDBSystem` object (`<database_name>.DBTypeForSAP.HDB.SystemHome.<sldsystemhome>`).

Note the following:

- The LMDB does not support fully qualified domain names (FQDNs). We therefore recommend specifying physical databases without domain names. The data supplier will remove domains before sending physical databases to SLD.
- If you don't configure a value for each tenant database in the `global.ini`, the value configured at the system level applies. If you don't configure a value at the system level, the value of the `SLDSYSTEMHOME` parameter in the default system profile (`hana/shared/<sid>/profile/DEFAULT.PFL`) is used; or if this value hasn't been set, the host name from the landscape properties file (`hana/shared/<sid>/lm_structure/landscapeVariables.properties`).
- Configure a physical database for either all tenants or none. If some tenants are configured and some are not, the fallback values mentioned above are not used for those tenants without a configured value. In this case, the data supplier does not deliver this tenant completely and an error is written to the error log.
- Changing the `sldsystemhome` parameter in the `global.ini` file does not require a restart, whereas changing the default profile does.

Configure the Dependency Chain

In a system replication landscape, communication between primary and secondary systems is based on internal host names. This ensures that each site can resolve the host name of other sites and that the replication chain can switch seamlessly in the event of a takeover. To model the dependency relationship between databases in the replication chain in SLD, the physical databases must be mapped to the internal host names of the systems to which they belong.

In the `system_landscape_hostname_resolution` section of the `global.ini` file of the system database, create an entry for all internal host names in the system replication landscape and as the value, set the physical database available on the host. We recommend you do this in all systems.

The dependency relationship between databases will be modeled in SLD with the association `SAP_DatabaseReplicationDependency`.

Note

For more information about mapping internal host names between primary and secondary systems, see the section on host name resolution for system replication.

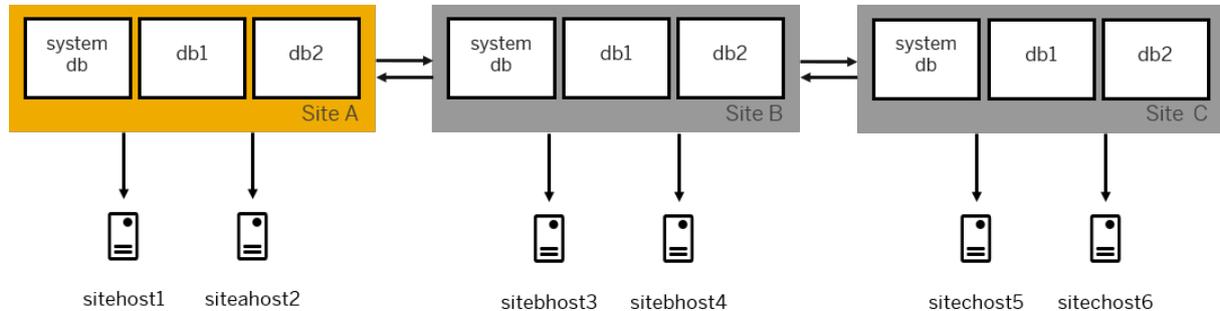
Association of Virtual and Physical Names (`enable_virtdbhome`)

Once the above configuration is in place, the virtual and physical databases must be associated. This takes place implicitly when the secondary system is registered with the primary system. With the `sr_register` command, the parameter `[sld] enable_virtdbhome` in the `nameserver.ini` file is set to **true** in the primary system and to **false** in the secondary system(s). This results in the virtual databases being associated with the physical databases of the primary system in SLD with the association `SAP_IdenticalDatabaseSystem`.

In the event of a takeover (`sr_takeover`), the value of this parameter is set to **true** in the new primary system and to **false** in the original primary system, thus changing the association of virtual databases to the physical databases in the new primary.

Example Configuration

The following figure shows a multitier system replication landscape. Each system in the chained setup is installed on two hosts for auto-host failover.



To send the required system landscape information to SLD, make the following landscape configuration settings in SAP HANA:

In all systems, configure virtual databases in the system database `global.ini` (section `system_landscape_hostname_virtualization`) as follows:

System	sldvirtdbhome Value for System	sldvirtdbhome Value for Tenant 1	sldvirtdbhome Value for Tenant 2
Primary system	systemDB	db1	db2
Tier 1 secondary system	systemDB	db1	db2
Tier 2 secondary system	systemDB	db1	db2

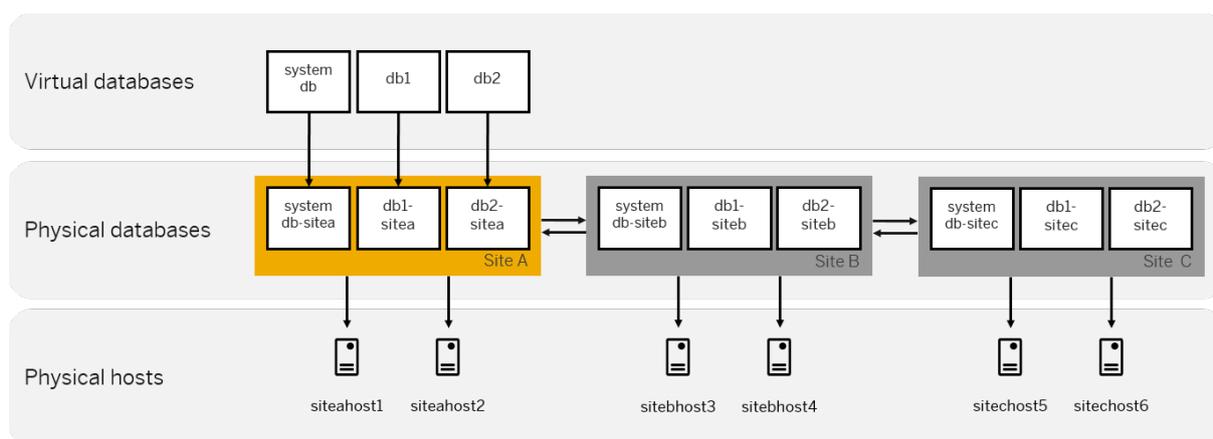
In all systems, configure the physical databases in the system database `global.ini` (section `system_landscape_hostname_virtualization`) as follows:

System	sldsystemhome Value for System	sldsystemhome Value for Tenant 1	sldsystemhome Value for Tenant 2
Primary system	systemDB-sitea	db1-sitea	db2-sitea
Tier 1 secondary system	systemDB-siteb	db1-siteb	db2-siteb
Tier 2 secondary system	systemDB-sitec	db1-sitec	db2-sitec

In all systems, configure the same host name resolution information in the system database `global.ini` (section `system_landscape_hostname_resolution`) as follows:

Parameter Name	Value for System	Value for Tenant 1	Value for Tenant 2
siteahost1	systemDB-sitea	db1-sitea	db2-sitea
siteahost2	systemDB-sitea	db1-sitea	db2-sitea

Parameter Name	Value for System	Value for Tenant 1	Value for Tenant 2
sitebhost3	systemDB-siteb	db1-siteb	db2-siteb
sitebhost4	systemDB-siteb	db1-siteb	db2-siteb
sitec host5	systemDB-sitec	db1-sitec	db2-sitec
sitec host6	systemDB-sitec	db1-sitec	db2-sitec



Related Information

[High Availability for SAP HANA \[page 1080\]](#)

[Host Name Resolution for System Replication \[page 1119\]](#)

[SLD Registration for SAP HANA SR Systems](#)

5.4.2 Central Monitoring and Administration with SAP Solution Manager

The monitoring and alerting infrastructure of SAP Solution Manager is based on a central agent infrastructure.

SAP delivers pre-configured agents for SAP HANA. If these agents have been deployed on SAP HANA and have been connected to SAP Solution Manager, SAP Solution Manager receives all alerts from the SAP HANA database. These alerts are displayed in the unified alert inbox.

SAP Solution Manager also provides an automated regular status check of your SAP solutions: SAP Early Watch Alert (EWA). This tool monitors the essential administrative areas of SAP components and keeps you up to date on their performance and stability. If you have any ABAP-based system in addition or running on SAP HANA, you can include the SAP HANA database information in the EWA report of this system: availability of services; size and growth; performance; important alerts; and correct setting of configuration parameters.

In addition to basic administration tasks, SAP provides specific task lists for SAP HANA, SAP Business Warehouse (BW) on SAP HANA, and SAP LT (Landscape Transformation) Replication Server. These lists

describe additional activities for administering these components and provide best practices for individual tasks.

5.4.3 Analyzing the Root Cause of Problems

You can use SAP Solution Manager to analyze the root cause of problems in your system landscape.

Often, SAP HANA systems are integrated with business applications that either call application accelerators in the SAP HANA database or are deployed on the database. If a problem occurs in this setup, you have to take a systematic approach to identify the precise cause of the fault. This standardized process must first broadly identify the components concerned and then analyze them more closely to arrive at the cause of the fault by process of elimination (top-down approach).

End-to-end root cause analysis in SAP Solution Manager provides your support organization with tools and methods that reliably identify the affected component while minimizing the time needed to solve the problem. In addition to your business applications, it supports also your SAP HANA database. This is the foundation for resolving problems in a holistic way. The DBA Cockpit is used in SAP Solution Manager to get a detailed insight into the status of the database. This is the same data that you can see in the SAP HANA cockpit for the SAP HANA database. But the DBA Cockpit supports other databases as well. If you have heterogeneous databases in your environment because your business applications still run on traditional databases, the DBA Cockpit enables you to use the same tool for the different databases.

Trace analysis records performance-related and functional faults in a specific user activity from the user's browser to the data stored in memory or on a storage sub-system. The measurement is triggered in the user interface and automatically activates recording of the relevant traces on every component processed by the user query.

After the root cause of a problem has been located, you can use detailed analysis to find a resolution. **Workload analysis** comprises server-related workload statistics for the connected systems. **Exception analysis** allows you to centrally analyze all exceptions from connected systems, such as serious error messages in logs or dumps. From here, you can start component-specific tools. And change analysis creates transparency for all changes (technical configuration, code, content) that have been made in the connected system landscape. This is particularly useful in the event of faults that occur once changes have been made to the production landscape.

Related Information

[DBA Cockpit for SAP HANA](#)

5.4.4 Controlling Change

In integrated system landscapes, SAP Solution Manager allows you to manage all changes centrally.

It is important to manage all changes in system landscapes using a central system. This is the only way to execute changes that affect more than one production system at the same time. For example, if you use SAP

HANA application-specific accelerators (for example, CO-PA) in combination with your SAP ERP back-end, you might need to synchronize changes of the data models on SAP HANA with changes for accessing the changed models from the transactions on SAP ERP. SAP Solution Manager provides central transport functions for the entire system landscape as well as additional support for quality management.

The process for transporting changes from a SAP HANA development system to a quality assurance and further on to the production system starts in the SAP HANA studio. There, you retrieve content from the SAP HANA source system and export it to enhanced Change and Transport System (CTS+) in SAP Solution Manager. The exported delivery units are attached automatically to a transport request of CTS+. After you have released this transport request, CTS+ triggers the automatic activation of the change in the SAP HANA repository on the target system (quality assurance or production).

In addition to supporting the transport and activation of new or changed data models, views, stored procedures, or analytic privileges, SAP Solution Manager offers additional possibilities to improve the quality of changes. You can use change analysis and reporting for getting information about the current status and history of changes. Configuration validation checks whether the systems of your system landscape (for example, development, quality assurance, and production) are configured consistently and according to the requirements. Quality Gate Management provides an additional quality inspection for projects and ensures changes are transported correctly and synchronously to the production systems. And Change Request Management within SAP Solution Manager controls the entire change execution with detailed process and workflow control. Out of these options, you can select those for controlling changes on SAP HANA that fit best to the quality strategy of your organization.

Related Information

[CTS+ How To Guides on SCN](#) 

6 System Administration

As a database administrator you are responsible for operational tasks related to the administration, monitoring, and maintenance of your SAP HANA systems.

Related Information

- [Aggregate Monitoring and Administration \[page 162\]](#)
- [Starting and Stopping SAP HANA Systems \[page 179\]](#)
- [Managing Tenant Databases \[page 189\]](#)
- [Configuring SAP HANA System Properties \(INI Files\) \[page 291\]](#)
- [Managing SAP HANA Licenses \[page 305\]](#)
- [Monitoring the SAP HANA Database \[page 317\]](#)
- [Managing and Monitoring the Performance of SAP HANA \[page 394\]](#)
- [Managing Tables \[page 479\]](#)
- [Workload Management \[page 621\]](#)
- [Scheduling of Recurring Administration Tasks \[page 658\]](#)
- [Getting Support \[page 659\]](#)

6.1 Aggregate Monitoring and Administration

Use the SAP HANA cockpit or SAP HANA studio to view high-level information about the status, availability, performance, capacity, and alert counts of all resources in your SAP HANA landscape before drilling down for database-level monitoring and administration.

Related Information

- [Managing Multiple Resources in SAP HANA Cockpit \[page 163\]](#)
- [Monitoring Multiple Systems in SAP HANA Studio \[page 176\]](#)

6.1.1 Managing Multiple Resources in SAP HANA Cockpit

Starting with an aggregate view of your registered resources in SAP HANA cockpit allows you to quickly discover any resources in your environment that have issues.

You can also pay attention to the groups of resources and decide whether you need to navigate into a group to investigate potential issues, or drill down to individual resources to obtain more details by clicking tabs, tiles, links, and numbers.

When you first access the cockpit, *My Resources*, the landscape level page, displays important high-level information about all the resources to which you have been granted access. A resource is an SAP HANA system, identified by a host and instance number, which may be a single- or multihost system. If you don't see any resources when you open the cockpit, either there are no resources registered in the cockpit, or your cockpit resource administrator has not assigned resources to you. See *Setting Up Cockpit with the Cockpit Manager* and *Working with Resources and Resource Groups*.

→ Tip

Selecting a resource name in the *Recently Accessed* list allows you to drill down to overview information for that individual resource. Unless your administrator has enabled single sign-on, you'll need to connect to the resource with a database user that has the system privilege CATALOG READ and SELECT on `_SYS_STATISTICS`.

Group Tabs

The tabs across the top of the page represent groups of resources. Each resource belongs to a usage type group (Production, Test, and Development) depending on configured system usage type. The resources may also belong to one or more groups created by the cockpit resource administrator. Included with each tab you see the overall number of high priority alerts for all resources in the group, as of the time of the most recent refresh. You can refresh the displayed data using the manual or auto-refresh icons in the top right corner.

Resource Directory

Selecting the *Resources* link beside the number of resources opens the *Resource Directory* which lists each resource, its connection, version, and resource groups to which it belongs. Selecting the name of any resources allows you to drill down to overview information for that individual resource.

You can use the cockpit to monitor and manage more than one resource, each running version SAP HANA 1.0 SPS 12, or later. Any resource running version SAP HANA 2.0 SPS 01, or later is set in multiple-container mode, by default. The cockpit can also monitor single-container systems running earlier versions of SAP HANA. When you drill down to the system *Overview*, and subsequently to *Manage Services*, the operations you have the option to perform depend on whether you have drilled down through the system database or the tenant. If you drill down to the overview for a system database you will see tenant statuses and alerts, as well as system database information. If you drill down to the overview for a tenant (or a single-container system), you will see only the statuses and alerts for that resource.

Alerts

[Top Resources with Alerts](#) shows you individual resources that require attention, based on the count of high priority alerts. Production systems are always displayed furthest to the left. Click the number representing the alert count for the resource to open the [Alerts Monitor](#) for the resource.

Additional Functionality

The links organized as [Administration](#), [Database Explorer](#), and [Help](#) each launch additional functionality. [Administration](#) links provide you with expanded capabilities for managing and monitoring groups of resources. If you are the cockpit administrator user, or a cockpit resource administrator user, the [Manage cockpit](#) link gives you access to the [Cockpit Manager](#). [Database Explorer](#) links allow developers and administrators to visually browse database objects (tables and schemas) and execute SQL statements.

Related Information

- [Set up SAP HANA Cockpit for the First Time \[page 43\]](#)
- [Determine Ports for SAP HANA Cockpit and Cockpit Manager \[page 45\]](#)
- [Open SAP HANA Cockpit \[page 47\]](#)
- [Open the Resource Directory \[page 165\]](#)
- [Using the Overview to Manage a Resource \[page 318\]](#)
- [Monitor Alerts from Multiple Resources \[page 169\]](#)
- [Monitor Aggregate Health \[page 170\]](#)
- [Comparing Configurations \[page 172\]](#)
- [Setup and Administration with the Cockpit Manager \[page 59\]](#)
- [Working with Resources and Resource Groups \[page 69\]](#)

6.1.1.1 Managing Groups of Resources

You can monitor and resolve issues by accessing information about a group of resources.

Each resource belongs to a usage type group (Production, Test, and Development) depending on configured system usage type. The resources may also belong to one or more groups created by the cockpit resource administrator.

On each you see the overall number of high priority alerts for all resources in the group, as of the time of the most recent refresh. (You can refresh the displayed data by using the manual or auto-refresh icons in the top right corner.) When you select a group, the cockpit displays information about resources in that specific group, rather than all the registered resources to which you have access.

Related Information

[Set up SAP HANA Cockpit for the First Time \[page 43\]](#)

[Determine Ports for SAP HANA Cockpit and Cockpit Manager \[page 45\]](#)

[Working with the Resource Directory \[page 165\]](#)

6.1.1.2 Working with the Resource Directory

The Resource Directory within SAP HANA cockpit contains information about all the registered resources belonging to resource groups to which you have been granted access.

For each resource, you can drill down for more information. Through the Resource Directory, you can also specify the database user credentials required to drill down to an individual resource, which is necessary unless single sign-on is in effect for that resource.

Related Information

[Open the Resource Directory \[page 165\]](#)

[Resource Details \[page 166\]](#)

[Resource Group Details \[page 166\]](#)

[Search, Sort, and Filter Tools for Resources and Groups \[page 167\]](#)

6.1.1.2.1 Open the Resource Directory

Display a list of registered resources to which you have been granted access.

Procedure

Related Information

[Resource Details \[page 166\]](#)

[Resource Group Details \[page 166\]](#)

[Search, Sort, and Filter Tools for Resources and Groups \[page 167\]](#)

[Connect to a Resource using Database Credentials \[page 168\]](#)

6.1.1.2.1.1 Resource Details

The Resource Directory provides these details about each registered resource:

Related Information

[Open the Resource Directory \[page 165\]](#)

[Resource Group Details \[page 166\]](#)

[Search, Sort, and Filter Tools for Resources and Groups \[page 167\]](#)

[Security Aspects of SAP HANA Cockpit \[page 86\]](#)

[Connect to a Resource using Database Credentials \[page 168\]](#)

6.1.1.2.1.2 Resource Group Details

On the Resource Groups tab of the Resource Directory in SAP HANA cockpit, you can see detailed information about resource groups.

The Resource Directory provides these details about each resource group:

Detail	Description
Name	Drill down on the resource group name to display the group's overview page.
Owner	The owner of the group. The owner shown for automatically generated groups (Production, Test, and Development) is "Unknown."
Description	The description provided by the group's creator.
Resources	The number of resources in this group.

Related Information

[Resource Details \[page 166\]](#)

[Search, Sort, and Filter Tools for Resources and Groups \[page 167\]](#)

[Managing Resource Groups \[page 80\]](#)

6.1.1.2.1.3 Search, Sort, and Filter Tools for Resources and Groups

Use the search, sort, and filter tools to work with lists of resources or resource groups.

Click [Resources](#) or [Resource Groups](#) at the top of the Resource Directory to choose your context.

Search	Enter a full or partial resource or group name in the Search box at the top of the screen and click the search icon  . The list is reduced to show only resources or groups that match your search string.
Change sorting rules	<p>You can modify the sorting rules to:</p> <ul style="list-style-type: none">• List resources or groups in ascending (the default) or descending order.• Sort resources by resource name (the default), connection, type, version, or resource group name.• Sort resource groups by resource group name (the default), owner, description, or resource name. <p>Click the settings icon , then click the sort icon . Choose sorting rules and click OK.</p>
Add a filter	<p>You can filter the list of resources by:</p> <ul style="list-style-type: none">• Type - select as many as you want. The options include HANA database, HANA system database, and HANA tenant database.• Version - select as many as you want. <p>Click the settings icon , then click the filter icon . Select the type of filter (version, for example) and select one or more options or click Select All. Click OK to save the filter.</p>
Modify a filter	<p>Each active filter appears in the blue bar at the top of the list of resources.</p> <p>To modify a filter, click the settings icon , click the filter icon , and select the type of filter. The filtering options appear. You can use the search tool to find the option you need.</p>
Cancel all filters	Click the settings icon  , and click the filter icon  . Click the clear filters icon  .
Cancel a single filter	Click the settings icon  , click the filter icon  , then select the type of filter to cancel. (The number of active filters of each type appears on the right.) Unselect the option for the filter you're cancelling and click OK .

Related Information

[Open the Resource Directory \[page 165\]](#)

[Resource Details \[page 166\]](#)

[Resource Group Details \[page 166\]](#)

6.1.1.2.2 Connect to a Resource using Database Credentials

Provide the credentials necessary to connect to a specific resource using SAP HANA cockpit.

Context

The cockpit resource administrator user may have used the Cockpit Manager configuration tool to enable cockpit to make use of the database's single sign-on (SSO) user authentication for a particular resource (running SAP HANA 2.0 SPS 01 or later). If not, each cockpit user needs to provide database user credentials in order to connect directly to the resource so as to drill down to information in the system [Overview](#). You can also connect using the `<sid>adm` user if you want the resource to be able to access the SAP Control process (which involves starting and stopping the resource, and restoring features).

→ Tip

Unless your administrator has enabled single sign-on, you'll need to connect to the resource with a database user that has the system privilege CATALOG READ and SELECT on `_SYS_STATISTICS`.

Procedure

1. In SAP HANA cockpit on the My Resources page or any Group Overview page, click [View resources directory](#) (under Administration).
2. At the top of the Resource Directory page, click [Resources](#) to display information about registered resources.
3. In the row displaying the resource to which you want to connect, choose to:
 - Select the link in the [SAPControl Credentials](#) column, and enter the name and password of the resource's `<sid>adm` user.
 - Select the link in the [Credentials](#) column. The wording on the link and in the subsequent dialog depend on how the cockpit resource administrator user has configured this resource in the Cockpit Manager:

Resource Configuration	Credentials Column displays...	Action
SSO has been enforced	SSO Enforced	You must connect to the resource through SSO. You don't need to enter database credentials.

Resource Configuration	Credentials Column displays...	Action
SSO has been allowed	<i>SSO Enabled</i> , or the most recently used database user name, with a link to <i>Choose Authentication</i>	Through the Choose Authentication dialog, you can choose to connect to the resource through SSO, or you can enter a different database user name and password.
SSO hasn't been allowed and you have not previously connected to the resource	The <i>Enter Credentials</i> link	Through the Enter Credentials dialog, you must enter a database user name and password. The cockpit securely stores and encrypts the credentials for next time.
SSO hasn't been allowed but you have previously connected to the resource	The most recently used database user name, with a link to <i>Manage Credentials</i>	Through the Manage Credentials dialog, you can choose to enter a database user name and password, or simply clear the previously used database credentials.

The cockpit securely encrypts and stores the credentials, and allows you to connect to the resource.

4. Select a resource name to drill down to that resource.

Related Information

[Resource Details \[page 166\]](#)

[Resource Group Details \[page 166\]](#)

[Search, Sort, and Filter Tools for Resources and Groups \[page 167\]](#)

6.1.1.3 Monitor Alerts from Multiple Resources

At a glance, you can see high-priority alerts from more than one resource.

Prerequisites

Your database user needs the object privilege SELECT on the schema `_SYS_STATISTICS`.

Context

In *My Resources*, and in the overview of any group, *Top Resources with Alerts* shows you individual resources that require attention, based on the count of high priority alerts. Production systems are always displayed furthest to the left. Clicking on the number representing the alert count for the resource opens the *Alerts Monitor* for the resource.

Procedure

1. Use [Top Resources with Alerts](#) to identify the alert you need to investigate.
2. Click on the alert to drill down to the [Alerts Monitor](#) for the resource that is showing the alert.
For more information about the [Alerts Monitor](#), see [Work with Alerts](#).

Related Information

[Managing Multiple Resources in SAP HANA Cockpit \[page 163\]](#)

[Monitoring Alerts \[page 341\]](#)

[Work with Alerts \[page 342\]](#)

6.1.1.4 Monitor Aggregate Health

Use the Aggregate Health Monitor to view high-level information about the running status, availability, performance, capacity, and alert counts of all your resources. You can drill down to see details about individual resources.

Prerequisites

Your database user needs the system privilege CATALOG READ.

Context

The Aggregate Health Monitor displays five high-level status indicators:

- Status - are managed resources running? Possible statuses are Running and, Stopped.
- Availability - are managed resources reachable on the network? Are they able to serve the business needs of their users, including humans and applications? Performance and capacity issues can affect availability.
- Performance - are managed resources meeting the response time expectations of database users, including humans and applications?
- Capacity - do managed resources have the system resources to support their applications?
- Alerts - do any managed resources need attention? Alert events, given priorities of high, medium, or low, are triggered when a resource exceeds state and range thresholds. The monitor displays the number of high and medium priority alerts.

Procedure

1. On the My Resources page or any resource group overview page, click *Monitor aggregate health* (under Administration).

If the list of resources is long, use the search, sort, and filter tools to find the resource you're interested in.

2. If a resource's alert counts are not 0/0, click the alert count numbers to drill down to the alerts page.
3. If any status indicator shows a problem, click the resource name to drill down to the overview for that resource.

Related Information

[Search, Sort, and Filter Tools in the Aggregate Health Monitor \[page 171\]](#)

6.1.1.4.1 Search, Sort, and Filter Tools in the Aggregate Health Monitor

Use the search, sort, and filter tools to work with the list of resources in the Aggregate Health Monitor.

Search

Enter a full or partial resource or group name in the search box at the top of the screen and click the search icon . The list is reduced to show only resources or groups that match your search string.

Change sorting rules

You can modify the sorting rules to:

- List resources in descending (the default) or ascending order
- Sort resources by issues score (the default), resource name, state, availability, performance, capacity, or alert count.

The issues score is a weighted ranking that identifies the resources most in need of attention. It's computed for each resource using all the indicators displayed in the Aggregate Health Monitor, with extra weight given (to produce a higher ranking) for resources that are not running (Status column) and resources that have high-priority alerts (Alert Counts column).

Click the settings icon , then click the sort icon . Choose sorting rules and click *OK*.

Add a filter

You can filter the list of resources by:

- Resource Type - select as many as you want. The options include HANA system, HANA MDB system, and HANA MDB tenant.
- Resource State - select as many as you want. The options include Running, Stopped, Error, Unknown, and Maintenance.
- Availability - select as many as you want. The options include High, Medium, Low, Normal, Not Collected, and Unsupported.
- Performance - select as many as you want. The options include High, Medium, Low, Normal, Not Collected, and Unsupported.
- Capacity - select as many as you want. The options include High, Medium, Low, Normal, Not Collected, and Unsupported.

Click the settings icon , then click the filter icon . Select the type of filter (resource type, for example) and select one or more options or click *Select All*. Click *OK* to save the filter.

Modify a filter

Each active filter appears in the blue bar at the top of the list of resources.

To modify a filter, click the settings icon , click the filter icon , and select the type of filter. The filtering options appear. You can use the search tool to find the option you need.

Cancel all filters

Click the settings icon  and click the filter icon . Click the clear filters icon .

Cancel a single filter

Click the settings icon , click the filter icon , then select the type of filter to cancel. (The number of active filters of each type appears on the right.) Unselect the option for the filter you're cancelling and click *OK*.

6.1.1.5 Comparing Configurations

Compare the configuration settings for two managed resources, capture snapshots of configurations, and compare snapshots.

Related Information

[Take a Snapshot of a Resource's Configuration \[page 173\]](#)

[Compare Resource Configurations \[page 174\]](#)

6.1.1.5.1 Take a Snapshot of a Resource's Configuration

Save a configuration snapshot: a timestamped copy of a managed resource's full set of configuration parameters.

Prerequisites

- Register the resource whose configuration you want to capture.

Context

Snapshots let you capture an accurate record of each resource's configuration; track configuration changes; and provide context to the historical data SAP HANA cockpit collects.

Procedure

1. Connect to the cockpit and sign in.

The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

The port number was configured during cockpit installation.

2. On any group overview page, click the *Compare configurations* link under Administration.
3. In the Configuration Manager's left pane:
 - Select the target resource's type, or
 - Select the target resource's group
4. Select the resource whose configuration you want to capture.

To sort the list of resources, click the Sort icon ↑↓ and choose sorting rules.

5. If you haven't authenticated with the resource you selected, sign in at the prompt.
6. Click the *Snapshots* tab in the middle of the screen.

The cockpit lists any previous snapshots of this resource's configuration.

7. Click *Take Snapshot*.
8. (Optional) Enter a description for the snapshot. The cockpit automatically associates the snapshot with its resource.
9. (Optional) To delete a snapshot, select it in the snapshots list and click *Delete Snapshot*.

6.1.1.5.2 Compare Resource Configurations

Compare the current configurations of two resources, compare two snapshots, or compare a current resource configuration to a snapshot.

Prerequisites

- Add the resources whose configuration you want to compare.
- Ensure that the resource or resources whose configurations you want to compare support configuration management. SAP HANA systems support configuration management in version SAP HANA 1.0 SPS 11 and later.

Context

To run a comparison, you select a source and a target to compare.

- When you compare two current configurations, the source and target resources must be of the same type (two SAP HANA systems, for example) and must be running the same software version.
- When you compare a current configuration to a snapshot, they must belong to the same system. That is, you cannot compare a current configuration to a snapshot of another system.
- When you compare two snapshots, they must belong to the same system. That is, you cannot compare a snapshot of one system to a snapshot of another system.

Procedure

1. Connect to the cockpit and sign in.

The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

The port number was configured during cockpit installation.

2. On any group overview page, click the [Compare configurations](#) link under Administration.
3. In the Configuration Manager's left pane:
 - Select the source resource's type, or
 - Select the source resource's group
4. Select the resource whose current or snapshot configuration you want to compare.

To sort the list of resources, click the Sort icon  and choose sorting rules.

5. If you haven't authenticated with the resource you selected, sign in at the prompt.

If you have authenticated but want to sign in as a different user, click [Authenticate](#) and enter new credentials.

6. To compare the current configurations of two resources:
 - a. With the source resource selected in the left pane, click *Compare* at the bottom of the screen on the left.
 - b. On the Configuration Comparison screen, select the target resource from the *Target Resource* drop-down list.

You can hover your mouse over an item in the *Target Resource* list to display connection details for that resource.

When you select a target resource, the cockpit displays comparison results for the parameters in the first group or configuration file (for example, `attributes.ini`). The group name (if applicable) or file name appears in the category drop-down list on the left, immediately above the parameter table.

- c. (Optional) To display in the results only parameters whose values are different in the source and the target, select the *Show differences only* checkbox.

The cockpit filters out of the results those parameters whose values are the same in the source and the target.

- d. To see comparison results for the remaining parameter groups or configuration files, pull down the category list on the left (immediately above the parameter table). Each parameter group or configuration file with differences is annotated with a count (for example, `attributes.ini (4)`). If you select *Show differences only* and the category list is empty, there are no differences between the two configurations.
 - e. To search on any string that appears in the currently displayed parameter table, enter a search term in the Search field and click the Search icon . The list is reduced to show only items that match your search term.

Click the blue  in the Search box to cancel the search and restore the full parameter table.

7. To compare a current configuration to a snapshot:
 - a. With the source resource selected in the left pane, click the *Parameters* tab in the middle of the screen.
 - b. Click *Compare* in the lower right corner of the screen.

Note

There are two *Compare* buttons at the bottom of the screen—be sure to click the one on the right. Its hover text says `Compare system configuration to a specific snapshot`.

- c. On the Configuration Comparison screen, select the target snapshot from the *Target Snapshot* drop-down list.

The Target Snapshot menu lists the snapshots by timestamp. Hover your mouse over the timestamp in the list to see that snapshot's description (if it has one).

The cockpit displays comparison results for the parameters in the first group or configuration file (for example, `attributes.ini`). The group name (if applicable) or file name appears in the drop-down menu on the left immediately above the parameter table.

- d. (Optional) To display in the results only parameters whose values are different in the source and the target, select the *Show differences only* checkbox.

The cockpit filters out of the results those parameters whose values are the same in the source and the target. If there are no differences between the source and the target in the current configuration file, the cockpit displays `No differences found` in the parameter table.

- e. To see comparison results for the remaining parameter groups or configuration files, pull down the category list on the left (immediately above the parameter table). Each parameter group or

configuration file with differences is annotated with a count (for example, `attributes.ini (4)`). If you select *Show differences only* and the category list is empty, there are no differences between the configuration and the snapshot.

- f. To search on any string that appears in the currently displayed parameter table, enter a search term in the Search field and click the Search icon . The list is reduced to show only items that match your search term.

Click the blue  in the Search box to cancel the search and restore the full parameter table.

8. To compare two snapshots:
 - a. With the resource selected in the left pane, click the *Snapshots* tab in the middle of the screen.
 - b. Select the source snapshot from the list.
 - c. Select *Compare Snapshots*.
 - d. On the Configuration Comparison screen, select the target snapshot from the *Target Snapshot* drop-down list.

The Target Snapshot list shows the snapshots by timestamp. Hover your mouse over the timestamp in the list to see that snapshot's description (if it has one).

The cockpit displays comparison results for the parameters in the first group or configuration file (for example, `attributes.ini`). The group name (if applicable) or file name appears in the category drop-down list on the left immediately above the parameter table.

- e. (Optional) To display in the results only parameters whose values are different in the source and the target, select the *Show differences only* checkbox.

The cockpit filters out of the results those parameters whose values are the same in the source and the target. If there are no differences between the source and the target in the currently displayed configuration settings, the cockpit displays `No differences found` in the parameter table.

- f. To see comparison results for the remaining parameter groups or configuration files, pull down the category list on the left (immediately above the parameter table). Each parameter group or configuration file with differences is annotated with a count (for example, `attributes.ini (4)`). If you select *Show differences only* and the category list is empty, there are no differences between the two snapshots.
- g. To search on any string that appears in the currently displayed parameter table, enter a search term in the Search field and click the Search icon . The list is reduced to show only items that match your search term.

Click the blue  in the Search box to cancel the search and restore the full parameter table.

6.1.2 Monitoring Multiple Systems in SAP HANA Studio

Using the *System Monitor* view of the SAP HANA studio, you can monitor the status of all registered systems.

Related Information

[System Monitor \[page 177\]](#)

[Options for Customizing the System Monitor \[page 178\]](#)

6.1.2.1 System Monitor

The *System Monitor* provides you with an overview of all your SAP HANA systems at a glance, including operational status, resource usage, and current alerts. From the *System Monitor*, you can drill down into the details of an individual system in the Administration editor.

Note

To see all information for all systems, you must have either the MONITORING role or the system privilege CATALOG READ and the object privilege SELECT on the schema _SYS_STATISTICS.

System ID	Operational State	Alerts	Data Disk (GB)	Log Disk (GB)	Trace Disk (GB)	Database Resident Memory (G...	System Resident Memory (G...	Used Memory (GB)	CPU (%)
ABS (DBA)	All services started	1 alert with HIGH priority, 1 alert wit...	525,74/4031,73	49,66/4031,73	0,21/4031,73	132,75/960,13	157,80/960,13	336,03/895,35	0
API (DBA)	All services started	1 alert with HIGH priority	2,80/787,43	2,30/787,43	0,11/787,43	12,71/94,48	42,79/94,48	38,25/87,15	3
GIS (DBA)	All services started	1 alert with HIGH priority, 1 alert wit...	64,84/3789,34	6,15/3789,34	0,17/3789,34	83,42/504,89	235,74/504,89	70,90/485,13	0
SHI (DBA)	All services started	1 alert with MEDIUM priority	3,84/914,92	2,30/914,92	0,02/914,92	9,48/23,45	17,20/23,45	18,34/21,10	14
UTD (DBA)	System status cannot be determined								
WA2 (DBA)	All services started	1 alert with MEDIUM priority	4,60/196,86	2,18/196,86	0,03/37,39	9,71/31,36	9,65/31,36	15,84/28,23	68

The System Monitor

Related Information

[Systems View \[page 121\]](#)

[Options for Customizing the System Monitor \[page 178\]](#)

Information Available in the System Monitor

The following information is available in the *System Monitor*.

Column	Description
System ID	ID assigned to system when added
Operational State	Overall system status
Alerts	The system issues alerts when resource usage and statistical thresholds are violated. These alerts are categorized as low, medium, or high priority. There are also information alerts. The number of alerts and their status is shown here.
Data Disk (GB)	Size of the data volume on disk
Log Disk (GB)	Size of the log volume on disk
Trace Disk (GB)	Size of trace files on disk
Database Resident Memory (GB)	Size of resident memory at operating system level owing to SAP HANA database processes
System Resident Memory (GB)	Total size of resident memory in the operating system
Used Memory (GB)	Amount of physical memory used by the SAP HANA database

Column	Description
CPU (%)	Percentage of CPU used by the SAP HANA database
Hostname	Name of the server hosting the SAP HANA database
Instance Number	Instance number is the administrative unit that comprises the server software components
System Data Disk (GB)	Total disk space occupied on disk(s) containing data
System Log Disk (GB)	Total disk space occupied on disk(s) containing log files
System Trace Disk (GB)	Total disk space occupied on disk(s) containing trace files
System Physical Memory (GB)	Total amount of physical memory used
System CPU (%)	Overall CPU usage
Distributed	Indicates whether the system is running on a single host or it is a distributed system running on more than one host
Start Time First	Time that the first service started This value is updated when system is restarted for any reason.
Start Time Latest	Time that the last service was started, if, for example, one of the services was re-started individually
Version	Software version number of the SAP HANA studio
Platform	Operating system on which the SAP HANA studio is running
Number of Crash Dump Files	The number of crash dump files in the trace directory of the system

6.1.2.2 Options for Customizing the System Monitor

The toolbar of the *System Monitor* provides options for customizing the view.

Toolbar Option	Description
 (Filter)	Allows you to select a sub-set of systems to display in the <i>System Monitor</i> , if for example you have a very large number of systems
 (Properties)	Allows you to configure properties of the <i>System Monitor</i> , such as the refresh interval and whether or not you want it to open automatically when you open the SAP HANA studio
 (Configure Viewer)	Allows you to configure which information is displayed, that is, which columns are visible

6.2 Starting and Stopping SAP HANA Systems

As the operating system administrator (<sid>adm user), you can stop, start, and restart an SAP HANA system using the SAP HANA cockpit or the SAP HANA studio. You can also stop and start a system from the command line using the SAPControl program.

Related Information

[Starting and Stopping Systems in SAP HANA Studio \[page 181\]](#)

[Starting and Stopping Systems in SAP HANA Cockpit \[page 179\]](#)

[Starting and Stopping Distributed SAP HANA Systems Using SAPControl \[page 1448\]](#)

6.2.1 Starting and Stopping Systems in SAP HANA Cockpit

You can use the SAP HANA cockpit to stop or start an SAP HANA system.

Related Information

[Open SAP HANA Cockpit \[page 47\]](#)

6.2.1.1 Start a Resource

Use SAP HANA cockpit to start a resource.

Prerequisites

You have the credentials of the operating system user (<sid>adm user) that was created when the system was installed.

Procedure

1. In SAP HANA cockpit, from the *My Resources* or Group Overview, drill down to the system *Overview*.

2. On the *Overall Database Status* tile (for a single container) or the *Overall System Database Status* tile (for a system database), click *Start System*.

Results

The database services start one by one, including those of any tenant databases. For details on starting an individual tenant database, see *Start a Tenant Database*.

When all services have started, the system has the status *Running*.

→ Tip

To analyze any problems that may occur during startup, you can access the system's diagnosis files from the homepage of the SAP HANA cockpit.

If you're unable to start a resource that was registered while it was unreachable, check the information entered during registration. The cockpit can't check the registration information for an unreachable resource, and thus can't tell the difference between a host or resource that's unreachable and one that doesn't exist. In particular, make sure these are correct:

- Host name
- Instance number
- Technical user name and password
- SAP HANA system ID

If you find an error, unregister the resource and register it again.

Related Information

[Start a Tenant Database \[page 213\]](#)

[Manage Services \[page 321\]](#)

[Register a Resource \[page 70\]](#)

[Unregister a Resource \[page 76\]](#)

6.2.1.2 Stop a Resource

Use SAP HANA cockpit to stop a resource.

Prerequisites

You have the credentials of the operating system user (<sid>adm user) that was created when the system was installed.

Procedure

1. In SAP HANA cockpit, from the [My Resources](#) or Group Overview, drill down to the system [Overview](#).
2. On the [Overall Database Status](#) tile (for a single container) or the [Overall System Database Status](#) tile (for a system database), click [Stop System](#).
3. Specify how you want to stop the system:

Option	Description
Softly	The system is stopped after all running statements have finished. If the system doesn't stop before the specified timeout, it is stopped immediately. The default timeout is 5 minutes.
Immediately	The system is stopped immediately. Open transactions are aborted and rolled back.

Results

The database services stop one by one. The services of tenant databases are stopped. For more information about how to stop an individual tenant database, see [Stop a Tenant Database](#) in the *SAP HANA Administration Guide*.

When all services have stopped, the system has the status [Stopped](#).

→ Tip

To analyze problems even when the system is stopped, you can access the system's diagnosis files from the homepage of the SAP HANA cockpit.

Related Information

[Manage Services \[page 321\]](#)

[Stop a Tenant Database \[page 214\]](#)

6.2.2 Starting and Stopping Systems in SAP HANA Studio

Use the SAP HANA studio to stop, start, or restart an SAP HANA system.

Related Information

[Start a System \[page 182\]](#)

[Stop a System \[page 183\]](#)

[Restart a System \[page 184\]](#)

[Stop and Start a Database Service \[page 186\]](#)

[Monitoring SAP HANA Systems During Stop and Start \[page 187\]](#)

[Restart Sequence \[page 188\]](#)

6.2.2.1 Start a System

Use the SAP HANA studio to start an SAP HANA system. All tenant databases will be started except those that were individually stopped.

Prerequisites

You have the credentials of the operating system user (<sid>adm user) that was created when the system was installed.

Procedure

1. In the *Systems* view, right-click the system you want to start and choose ► *Configuration and Monitoring* ► *Start System...* ►

i Note

Execute the start command from the system database. The *Start System...* command is not available from tenant databases. For more information about how to stop an individual tenant database, see *Stop and Start a Tenant Database*.

2. Enter the user name and password of the operating system administrator that was created when the system was installed (that is, <sid>adm user).

Results

- The Administration editor opens in diagnosis mode and the database services start one by one. When all services have started, the system appears as operational () in the *Systems* view.
- All tenant databases are started. However, if a tenant database was previously stopped individually, it is not started with the system. For more information about how to stop an individual tenant database, see *Stop and Start a Tenant Database*.

→ Tip

Refresh the *Systems* view to update the status of other instances of the system or tenant databases registered in the SAP HANA studio.

For more information about starting a distributed SAP HANA system using the `sapcontrol` program, see *Starting and Stopping a Distributed SAP HANA System Using sapcontrol*.

Related Information

[Operating System User sidadm \[page 714\]](#)

[Start a Tenant Database \[page 213\]](#)

[Monitoring System Availability \[page 360\]](#)

[Starting and Stopping Distributed SAP HANA Systems Using SAPControl \[page 1448\]](#)

6.2.2.2 Stop a System

In certain situations, you may have to stop your system, for example after changing certain system parameters. Use the SAP HANA studio to stop an SAP HANA system.

Prerequisites

You have the credentials of the operating system user (<sid>adm user) that was created when the system was installed.

Procedure

1. In the *Systems* view, right-click the system you want to stop and choose ► *Configuration and Monitoring* ► *Stop System...* ►

i Note

Execute the stop command from the system database. The *Stop System...* command is not available from tenant databases. For more information about how to stop an individual tenant database, see *Stop and Start a Tenant Database*.

2. Specify how you want to stop the system:

Option	Description
Soft	The system is stopped after all running statements have finished or the specified timeout is reached.
Hard	The system is stopped immediately. Open transactions are aborted and rolled back.

3. Optional: Specify a stop wait timeout (date and time).
If the system does not shut down before the specified timeout, it is shut down forcefully.

4. Enter the user name and password of the operating system user that was created when the system was installed (that is, `<sid>adm` user).

Results

- The Administration editor opens in diagnosis mode and the database services stop one by one. When all services have stopped, the system appears as non-operational () in the *Systems* view.
- All tenant databases are stopped. For more information about how to stop an individual tenant database, see *Stop and Start a Tenant Database*.

→ Tip

Refresh the *Systems* view to update the status of other instances of the system or tenant databases registered in the SAP HANA studio.

For more information about stopping a distributed SAP HANA system using the `sapcontrol` program, see *Starting and Stopping a Distributed SAP HANA System with sapcontrol*.

Related Information

[Operating System User sidadm \[page 714\]](#)

[Monitoring System Availability \[page 360\]](#)

[Start a Tenant Database \[page 213\]](#)

[Starting and Stopping Distributed SAP HANA Systems Using SAPControl \[page 1448\]](#)

6.2.2.3 Restart a System

In certain situations, you may have to restart the system, for example, after a power failure. Use the SAP HANA studio to restart an SAP HANA system.

Prerequisites

You have the credentials of the operating system user (`<sid>adm` user) that was created when the system was installed.

Procedure

1. In the *Systems* view, right-click the system you want to start and choose ► *Configuration and Monitoring* ► *Restart System...* ►

i Note

Execute the restart command from the system database. The *Restart System...* command is not available from tenant databases. For more information about how to stop a tenant database, see *Stop and Start a Tenant Database*.

2. Specify how you want to stop the system:

Option	Description
Soft	The system is stopped after all running statements have finished or the specified timeout is reached.
Hard	The system is stopped immediately. Open transactions are aborted and rolled back.

3. Optional: Specify a stop wait timeout (date and time).
4. Enter the user name and password of the operating system user that was created when the system was installed (that is, <sid>adm user).

Results

- The Administration editor opens in diagnosis mode. The database services first stop one by one and then restart one by one. The icon displayed for the system in the *Systems* view changes as the status of the services changes.
- All tenant databases are stopped and restarted. However, if a tenant database was previously stopped individually, it is not restarted with the system. For more information about how to stop an individual tenant database, see *Stop and Start a Tenant Database*.

→ Tip

Refresh the *Systems* view to update the status of other instances of the system or tenant databases registered in the SAP HANA studio.

Related Information

[Restart Sequence \[page 188\]](#)

[Operating System User sidadm \[page 714\]](#)

[Start a Tenant Database \[page 213\]](#)

6.2.2.4 Stop and Start a Database Service

You can stop and start individual database services (`nameserver`, `indexserver`, `xsengine` and so on) running on an SAP HANA host or hosts.

Prerequisites

You have the system privilege `SERVICE ADMIN`.

Context

You may need to stop and (re)start services in the following situations, for example:

- A host in a distributed system failed and a standby host took over. However, the services of the failed host remain inactive even after the host is reachable again. In this case, you need to restart the services manually.
- After an update of SAP HANA extended application services (SAP HANA XS), the `xsengine` service needs to be restarted.

i Note

The SAP HANA database provides several features in support of high availability, one of which is service auto-restart. In the event of a failure or an intentional intervention by an administrator that disables one of the SAP HANA services, the SAP HANA service auto-restart function automatically detects the failure and restarts the stopped service process.

Procedure

1. In the Administration editor open the **► Landscape ► Services ►** tab.
2. Right-click the service and choose the required option:

Option	Description
Stop...	The service is stopped normally and then typically restarted.
Kill...	The service is stopped immediately and then typically restarted To have the system create a crash dump file, select the option <i>Create Core File</i> . You can access the generated crash dump file on the <i>Diagnosis Files</i> tab.
Reconfigure Service...	The service is reconfigured. This means that any changes made to parameters in the system's configuration files are applied.
Start Missing Services...	Any inactive services are started.

Related Information

[Monitoring Status and Resource Usage of System Components \[page 364\]](#)

[High Availability for SAP HANA \[page 1080\]](#)

[View Diagnosis Files in SAP HANA Studio \[page 663\]](#)

6.2.2.5 Monitoring SAP HANA Systems During Stop and Start

You can access the diagnosis files of a system that is starting up or has stopped by opening the Administration editor in the diagnosis mode.

The SAP HANA studio normally collects information about the system using SQL. However, when the system has not yet started, no SQL connection is available. Therefore, while the system is starting up or is stopped, the SAP HANA studio collects information about the database using the connection of the SAP start service (`sapstartsrv`). If you have the credentials of the operating system administrator (user `<sid>adm`), you can view this information in the Administration editor in diagnosis mode.

In this way, you can analyze any problems that may occur during startup or while the system is stopped. You can also read diagnosis files even when the system is stopped.

The Administration editor opens automatically in diagnosis mode in the following situations:

- When you open the Administration editor for a system without an SQL connection
- When you initiate the start, stop, or restart of a system

You can manually open a system in diagnosis mode by choosing the  (*Open Diagnosis Mode*) button from the drop-down menu of the  (*Administration*) button in the *Systems* view.

Related Information

[Operating System User sidadm \[page 714\]](#)

[Monitoring Overall System Status and Resource Usage \[page 362\]](#)

[Troubleshooting an Inaccessible or Unresponsive SAP HANA System \[page 684\]](#)

6.2.3 Starting and Stopping Systems with SAPControl

You can use the SAPControl program to start or stop SAP HANA system from the command line.

i Note

You must be logged on to the SAP system host as user `<sid>adm` or as a user with root permissions.

Action	Command
Start the system	<code>/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StartSystem HDB</code>
Stop the system	<code>/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StopSystem HDB</code>
Query current status of all hosts in the system	<code>/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function GetSystemInstanceList</code>

6.2.4 Restart Sequence

The SAP HANA system restart sequence restores the system to a fully operational state quickly.

When you restart an SAP HANA system, the following activities are executed by the restart agent of the persistence layer.

1. The data volume of each service is accessed in order to read and load the restart record.
2. The list of open transactions is read into memory.
3. Row tables are loaded into memory.
4. Open transactions are processed using the redo log:
 1. Write transactions that were open when the database was stopped are rolled back.
 2. Changes of committed transactions that were not written to the data area are rolled forward.

The first column tables start being reloaded into memory as they are accessed for roll forward.

i Note

Since a regular or "soft" shutdown writes a savepoint, there are no replay log entries to be processed in this case.

After this step, the database is technically available and logon is possible.

5. Aborted transactions are determined and rolled back.
6. A savepoint is performed with the restored consistent state of the database.
7. Column tables that are marked for preload and their attributes are asynchronously loaded in the background (if they have not already been loaded as part of log replay).
The preload parameter is configured in the metadata of the table. This feature is useful for example to make certain tables and columns used by important business processes available more quickly.
8. Column tables that were loaded before restart and their attributes start reloading asynchronously in the background (if they have not already been loaded as part of log replay or because they are marked for preload).
During normal operation, the system tracks the tables currently in use. This list is used as basis for reloading tables after a restart.

Reloading column tables as described in steps 7 and 8 restores the database to a fully operational state more quickly. However, it does create performance overhead and may not be necessary in non-production systems. You can deactivate the reload feature in the `indexserver.ini` file by setting the `reload_tables` parameter in the `sql` section to **false**. In addition, you can configure the number of tables whose attributes are loaded in parallel using the `tables_preloaded_in_parallel` parameter in the `parallel` section of `indexserver.ini`. This parameter also determines the number of tables that are preloaded in parallel.

6.3 Managing Tenant Databases

As the administrator of a tenant database system, you are responsible for creating and configuring new tenant databases, subsequently monitoring the availability and performance of databases, as well as performing certain database administration tasks.

i Note

Administration of tenant databases is possible using the SAP HANA cockpit. However, command-line tools are required for some tasks.

i Note

If you have SAP HANA options installed, review the section about tenant databases in the administration guide of the corresponding option for additional information before proceeding. Be aware that you need additional licenses for SAP HANA options and capabilities. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities*.

Related Information

[Creating and Configuring Tenant Databases \[page 189\]](#)

[Monitoring and Managing Tenant Databases \[page 243\]](#)

[Configuring Memory and CPU Usage for Tenant Databases \[page 266\]](#)

[Important Disclaimer for Features in SAP HANA Platform \[page 1980\]](#)

6.3.1 Creating and Configuring Tenant Databases

You create tenant databases after installation if no initial tenant was created, after conversion from a single-container system to a multiple-container system, or anytime a new database is needed.

As a system administrator, you create tenant databases from the system database. You can then configure the new databases as required:

- Increase the database isolation level
- Disable certain features that are not required in tenant databases (for example, backup operations). Disabled features in tenant databases can still be accessed through the system database.
- Enable and configure cross-database access if read-only queries between tenant databases is required
- Edit the configuration change blacklist so that critical system properties cannot be changed by tenant database administrators
- Configure the SAP Web Dispatcher if tenant databases will be accessed by HTTP clients via the SAP HANA XS classic server

i Note

Administration of tenant databases is possible using the SAP HANA cockpit. However, command-line tools are required for some tasks.

Related Information

[Converting an SAP HANA System to Support Tenant Databases \[page 190\]](#)

[Increase the System Isolation Level \[page 202\]](#)

[Create a Tenant Database \[page 210\]](#)

[Disable Features on a Tenant Database \[page 217\]](#)

[Enable and Configure Cross-Database Access \[page 219\]](#)

[Prevent Changes to System Properties in Tenant Databases \[page 225\]](#)

[Configure HTTP\(S\) Access to Tenant Databases via SAP HANA XS Classic \[page 1578\]](#)

[Administration of Tenant Databases \[page 20\]](#)

6.3.1.1 Converting an SAP HANA System to Support Tenant Databases

You can convert an SAP HANA system to support tenant databases using the SAP HANA database lifecycle manager (HDBLCM) resident program. Converting an SAP HANA system to a tenant database system is permanent and cannot be reversed.

If your system was installed in single-container mode, you can still implement tenant databases by converting the system to a tenant database system. During the conversion process, the system database and one tenant database are created. The tenant database contains all the data of the original system, including users, system configuration, connection properties (port configuration), and system license. However, it does **not** contain the backup history.

After conversion, you can create and configure further tenant databases as needed.

i Note

After conversion, a port offset value of 100 is used to reserve ports for system replication communication. A port offset that you defined before the conversion is not changed.

Related Information

[Convert to Tenant Databases Using the Graphical User Interface \[page 191\]](#)

[Convert to Tenant Databases Using the Command-Line Interface \[page 193\]](#)

[Convert to Tenant Databases Using the Web User Interface \[page 195\]](#)

6.3.1.1.1 Convert to Tenant Databases Using the Graphical User Interface

You can convert an SAP HANA system to support tenant databases using the SAP HANA database lifecycle manager (HDBLCM) resident program in the graphical user interface. Converting an SAP HANA system to a tenant database system is permanent and cannot be reversed.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).
- The host has access to the installation directories `<sapmnt>` and `<sapmnt>/<SID>`.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on as root user or as the system administrator user `<sid>adm`.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblcml
```

By default, `<sapmnt>` is `/hana/shared`.

2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

```
./hdblcmlgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

3. Select *Convert to Tenant Databases* from the activity options. Then select *Next*.
4. Provide the password of the `<sid>adm` user and the `SYSTEM` user of `SYSTEMDB`, then select *Next*.
5. Review the summary, and select *Run* to finalize the configuration.

Results

Your SAP HANA system is a tenant database system with one system database and one tenant database, both of which are running. You can verify this by adding both databases to SAP HANA cockpit and querying the public view `M_DATABASES` from the system database. The result will look like this:

DATABASE_NAME	DESCRIPTION	ACTIVE_STATUS
SYSTEMDB	SystemDB-<SID>-<INSTANCE>	YES
<SID>	SingleDB-<SID>-<INSTANCE>	YES

Note the following about the tenant database:

- It contains all the data (including users, configuration, and connection properties) of the original system (but not the original backup history).
- Configuration files that are tenant-specific (e.g. `indexserver.ini`, `xsengine.ini`, etc.) are now stored at the following location: `/usr/sap/<SID>/SYS/global/hdb/custom/config/DB_<database_name>`.
- Its trace files are now stored at the following location: `/usr/sap/<SID>/HDB<instance>/<host>/trace/DB_<database_name>`.

i Note

Any trace files that were in the trace directory before the system was converted are not moved.

Next Steps

- Create and configure any additionally required tenant databases. For more information, see *Create a Tenant Database*.

i Note

If you configured the properties of the index server, script server, or xsengine server in your original system, these settings initially apply to **all** new tenant databases. You must explicitly configure tenant database if required. For more information, see *System Properties in Tenant Database Systems* in the *SAP HANA Administration Guide*.

- If HTTP access via the SAP HANA XS classic server is required, update the configuration of the Web Dispatcher. For more information, see *Configure HTTP Access to Tenant Databases* in the *SAP HANA Administration Guide*.

Related Information

[Password Policy Configuration Options \[page 723\]](#)

[Create a Tenant Database \[page 210\]](#)

[Deploy a Delivery Unit Archive \(*.tgz\) \[page 996\]](#)

[Install a Permanent License \[page 315\]](#)

[Creating Backups \[page 1313\]](#)

[Database-Specific Configuration Parameters \[page 293\]](#)

[Configure HTTP\(S\) Access to Tenant Databases via SAP HANA XS Classic \[page 1578\]](#)

6.3.1.1.2 Convert to Tenant Databases Using the Command-Line Interface

You can convert an SAP HANA system to support tenant databases using the SAP HANA database lifecycle manager (HDBLCM) resident program in the command-line interface. Converting an SAP HANA system to a tenant database system is permanent and cannot be reversed.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).
- The host has access to the installation directories `<sapmnt>` and `<sapmnt>/<SID>`.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on as root user or as the system administrator user `<sid>adm`.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblcml
```

By default, `<sapmnt>` is `/hana/shared`.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblcml --action=convert_to_multidb
```

3. Provide the password of the `<sid>adm` user and `SYSTEM` user of `SYSTEMDB` user.
4. Review the summary, and select `y` to finalize the configuration.

Results

Your SAP HANA system is a tenant database system with one system database and one tenant database, both of which are running. You can verify this by adding both databases to SAP HANA cockpit and querying the public view `M_DATABASES` from the system database. The result will look like this:

DATABASE_NAME	DESCRIPTION	ACTIVE_STATUS
SYSTEMDB	SystemDB-<SID>-<INSTANCE>	YES
<SID>	SingleDB-<SID>-<INSTANCE>	YES

Note the following about the tenant database:

- It contains all the data (including users, configuration, and connection properties) of the original system (but not the original backup history).
- Configuration files that are tenant-specific (e.g. `indexserver.ini`, `xsengine.ini`, etc.) are now stored at the following location: `/usr/sap/<SID>/SYS/global/hdb/custom/config/DB_<database_name>`.
- Its trace files are now stored at the following location: `/usr/sap/<SID>/HDB<instance>/<host>/trace/DB_<database_name>`.

i Note

Any trace files that were in the trace directory before the system was converted are not moved.

Next Steps

- Create and configure any additionally required tenant databases. For more information, see *Create a Tenant Database*.

i Note

If you configured the properties of the index server, script server, or xsengine server in your original system, these settings initially apply to **all** new tenant databases. You must explicitly configure tenant database if required. For more information, see *System Properties in Tenant Database Systems* in the *SAP HANA Administration Guide*.

- If HTTP access via the SAP HANA XS classic server is required, update the configuration of the Web Dispatcher. For more information, see *Configure HTTP Access to Tenant Databases* in the *SAP HANA Administration Guide*.

Related Information

[Password Policy Configuration Options \[page 723\]](#)

[Create a Tenant Database \[page 210\]](#)

[Deploy a Delivery Unit Archive \(*.tgz\) \[page 996\]](#)

[Install a Permanent License \[page 315\]](#)

[Creating Backups \[page 1313\]](#)

[Database-Specific Configuration Parameters \[page 293\]](#)

[Configure HTTP\(S\) Access to Tenant Databases via SAP HANA XS Classic \[page 1578\]](#)

[import_content \[page 198\]](#)

[nostart \[page 198\]](#)

[nostart_tenant_db \[page 199\]](#)

6.3.1.1.3 Convert to Tenant Databases Using the Web User Interface

You can convert an SAP HANA system to support tenant databases using the SAP HANA database lifecycle manager Web user interface. Converting an SAP HANA system to a tenant database system is permanent and cannot be reversed.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).
- The host has access to the installation directories `<sapmnt>` and `<sapmnt>/<SID>`.
- The SAP HANA system has been installed or updated with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
 - Microsoft Windows
 - Internet Explorer - Version 9 or higher
If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing **Tools > Compatibility View Settings**.
 - Microsoft Edge
 - Mozilla Firefox - Latest version and Extended Support Release
 - Google Chrome - Latest version
 - SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
 - Mac OS - Safari 5.1 or higher

Note

For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html

i Note

The URL is case sensitive. Make sure you enter upper and lower case letters correctly.

SAP HANA cockpit	<ol style="list-style-type: none"> 1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. https://<host_FQDN>:<port>
-------------------------	--

i Note

FQDN = fully qualified domain name

2. Drill down on the name of the system from *My Resources* or from a group.
3. The links in *Platform Lifecycle Management* each launch additional functionality, giving you expanded capabilities for managing the resource.

2. Select the *Convert to Tenant Databases* tile.
3. Optional: Modify the following parameters in the *Advanced Parameters Configuration* dialog. To access the *Advanced Parameters Configuration* dialog, click on the gear icon in the footer bar of the SAP HANA HDBLCM Web user interface.

Option	Description
Import Delivery Units In The System Database	Import Delivery Units In The System Database
Do Not Start Instance After Reconfiguration	Do Not Start Instance After Reconfiguration
Do Not Start Tenant Database After Reconfiguration	Do Not Start Tenant Database After Reconfiguration
Timeouts	Sets customized timeouts (start_instance, stop_instance).

4. Provide the password of the <sid>adm user and SYSTEM user of SYSTEMDB user, then select *Next*.
5. Review the summary, and select *Run* to finalize the configuration.

Results

Your SAP HANA system is a tenant database system with one system database and one tenant database, both of which are running. You can verify this by adding both databases to SAP HANA cockpit and querying the public view M_DATABASES from the system database. The result will look like this:

DATABASE_NAME	DESCRIPTION	ACTIVE_STATUS
SYSTEMDB	SystemDB-<SID>--<INSTANCE>	YES

<SID>	SingleDB-<SID>-<INSTANCE>	YES
-------	---------------------------	-----

Note the following about the tenant database:

- It contains all the data (including users, configuration, and connection properties) of the original system (but not the original backup history).
- Configuration files that are tenant-specific (e.g. indexserver.ini, xsengine.ini, etc.) are now stored at the following location: `/usr/sap/<SID>/SYS/global/hdb/custom/config/DB_<database_name>`.
- Its trace files are now stored at the following location: `/usr/sap/<SID>/HDB<instance>/<host>/trace/DB_<database_name>`.

i Note

Any trace files that were in the trace directory before the system was converted are not moved.

Next Steps

- Create and configure any additionally required tenant databases. For more information, see *Create a Tenant Database*.

i Note

If you configured the properties of the index server, script server, or xsengine server in your original system, these settings initially apply to **all** new tenant databases. You must explicitly configure tenant database if required. For more information, see *System Properties in Tenant Database Systems* in the *SAP HANA Administration Guide*.

- If HTTP access via the SAP HANA XS classic server is required, update the configuration of the Web Dispatcher. For more information, see *Configure HTTP Access to Tenant Databases* in the *SAP HANA Administration Guide*.

Related Information

[SAPUI5 Developer Guide](#)

[Password Policy Configuration Options \[page 723\]](#)

[Create a Tenant Database \[page 210\]](#)

[Deploy a Delivery Unit Archive \(*.tgz\) \[page 996\]](#)

[Install a Permanent License \[page 315\]](#)

[Creating Backups \[page 1313\]](#)

[Database-Specific Configuration Parameters \[page 293\]](#)

[Configure HTTP\(S\) Access to Tenant Databases via SAP HANA XS Classic \[page 1578\]](#)

6.3.1.1.4 Parameter Reference: Converting an SAP HANA System to Support Tenant Databases

Parameters can be specified when converting an SAP HANA system to tenant databases in order to customize the configuration task.

The SAP HANA database lifecycle manager convert to multtidb action also supports the following parameters:

- batch
- configfile
- dump_configfile_template
- help
- list_systems
- read_password_from_stdin
- version

For more information about these parameters, see the *SAP HANA Server Installation and Update Guide*

For a complete list of the parameters, call the help of the convert to multtidb task with the following command:

```
./hdblcm --action=convert_to_multtidb --help
```

6.3.1.1.4.1 import_content

Imports delivery units.

Syntax

In the command line, the following syntax is used:

```
./hdbinst --import_content [=off]
```

Remarks

The default for this parameter is `--import_content`.

Related Information

[SAP HANA Content \[page 994\]](#)

[General Troubleshooting for the SAP HANA Platform LCM Tools \[page 946\]](#)

6.3.1.1.4.2 nostart

Prevents the SAP HANA system from being started.

Syntax

In the command line, the following syntax is used:

```
--nostart
```

6.3.1.1.4.3 `nostart_tenant_db`

Prevents the SAP HANA tenant databases from being started.

Syntax

In the command line, the following syntax is used:

```
--nostart_tenant_db
```

6.3.1.1.5 Convert a System Replication Landscape to Support Tenant Databases

An SAP HANA system replication landscape can be converted to support tenant databases in an offline or near-zero downtime approach.

You can convert an SAP HANA single database container system that is part of a system replication configuration to support tenant databases. You have the choice to convert your SAP HANA system after taking all sites offline, or to convert in a near-zero downtime approach to minimize system downtime.

Offline Conversion

To perform an offline conversion all systems at all sites must first be stopped. Then, starting with the primary site, execute the conversion script. Once the primary site is converted and online again, continue the conversion procedure with the next site, following the replication chain.

Near-Zero Downtime Conversion

In order to carry out a near-zero downtime conversion, execute the conversion script on the secondary site. Once the conversion of this site is complete, perform a takeover. After stopping all other systems, execute the conversion script on the primary site and reregister the primary system.

Related Information

[Perform an Offline Conversion \[page 200\]](#)

6.3.1.1.5.1 Perform an Offline Conversion

You can perform an offline conversion of an SAP HANA system replication landscape to support tenant databases. Converting an SAP HANA system to a tenant database system is permanent and cannot be reversed.

Prerequisites

- The statistics server is **not** running as a separate server process (`statisticsserver`), but instead as an embedded service in the master index server. If this is not the case, migrate the statistics server to the embedded statistics service as described in SAP Note 1917938.
- The SAP HANA system has been installed with its server software on a shared file system (export options `rw,no_root_squash`).
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- You are logged on as the system administrator user `<sid>adm`.

Procedure

1. Stop all SAP HANA systems on all sites.
2. Run the following command on the master host of the primary system:

```
<sapmnt>/<SID>/hdb1cm/hdb1cm --action=convert_to_multidb
```

By default, `<sapmnt>` is `/hana/shared`.

3. Specify a new system user password.
4. Wait until the conversion has finished and the system is active again.
5. Create a data backup of the system database.
6. Repeat steps 2 through 4 on all remaining secondary systems, following the replication chain.

Related Information

[SAP Note 1917938](#)

[Setting Up SAP HANA System Replication \[page 1098\]](#)

6.3.1.1.5.2 Perform a Near-Zero Downtime Conversion

You can perform a near-zero downtime conversion of an SAP HANA system replication landscape to support tenant databases. Converting an SAP HANA system to a tenant database system is permanent and cannot be reversed.

Prerequisites

- The statistics server is **not** running as a separate server process (`statisticsserver`), but instead as an embedded service in the master index server. If this is not the case, migrate the statistics server to the embedded statistics service as described in SAP Note 1917938.
- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- You are logged on as the system administrator user `<sid>adm`.

Procedure

1. Run the following command on master host of the secondary system:

```
<sapmnt>/<SID>/hdb lcm/hdb lcm --action=convert_to_multidb
```

By default, `<sapmnt>` is `/hana/shared`.

2. Specify a new system user password.
3. Wait until the conversion has finished and all systems are in sync again.
4. Stop all systems except the one you just converted.
5. Perform a takeover by the converted system.
6. Run the following command on master host of the primary system:

```
<sapmnt>/<SID>/hdb lcm/hdb lcm --action=convert_to_multidb --nostart
```
7. Reregister the original primary system as the new secondary system.

Related Information

[SAP Note 1917938](#)

[Setting Up SAP HANA System Replication \[page 1098\]](#)

6.3.1.2 Increase the System Isolation Level

You can increase the isolation level of an existing system from low (default) to high. With high isolation, the processes of individual tenant databases run under dedicated operating system (OS) users belonging to dedicated (OS) groups and internal communication is secured.

Prerequisites

- You have root access to the SAP HANA system.
- You are logged on to the system database in the SAP HANA cockpit.
- You have the system privilege `DATABASE ADMIN`.
- Internal SAP HANA communication has been appropriately configured for TLS/SSL.
`[communication] ssl` in the `global.ini` file must have the value `false` (default) or `systemPKI`.

⚠ Caution

If you are using a manually configured public key infrastructure (PKI) to secure internal communication between hosts, the property `[communication] ssl` must be `true`. You can switch to system PKI by changing setting the parameter to `systemPKI`.

For more information, see *Secure Internal Communication and Server-Side TLS/SSL Configuration Properties for Internal Communication* in the *SAP HANA Security Guide*.

- If the system is running in an SAP HANA system replication configuration, the system PKI SSFS data file and key file have been copied from the primary system to the same location on the secondary system(s):
 - `$DIR_INSTANCE/./global/security/rsecssfs/data/SSFS_<SID>.DAT`
 - `$DIR_INSTANCE/./global/security/rsecssfs/key/SSFS_<SID>.KEY`

Procedure

1. For every tenant database, create a dedicated OS user and group:
 - a. As root user, log on to the server on which the name server of the system database is running.
 - b. Create new groups for every tenant database:

```
groupadd <groupname>
```
 - c. Create new users for every tenant database, specifying `sapsys` as the primary group:

```
useradd -g sapsys <username>
```
 - d. Add every new user to the `sidshm` group and their own group as secondary groups:

```
usermod -G <sid>shm,<usergroup> <username>
```

i Note

If the system is distributed across multiple hosts, you must create identical users and groups on every host. Users and groups must have the same names and IDs on all hosts.

2. Stop all tenant databases in the system.

In the system database, execute the SQL statement `ALTER SYSTEM STOP DATABASE <database_name>`.

→ Tip

You can also stop tenant databases in the *Manage Databases* app of the SAP HANA cockpit.

3. Configure the system for high isolation.

As the operating system user `<sid>adm`, log on to the server on which the master index server is running and run the following command:

```
python /usr/sap/<SID>/HDB<instance>/exe/python_support/convertMDC.py --
change=databaseIsolation --isolation=high
```

This command runs the following actions:

- Stops the system
 - Changes the value of the `[multidb] database_isolation` property in the `global.ini` file to `high`
 - Starts the system
4. Assign every database to their respective OS user and group.

In the system database, execute the SQL statement `ALTER DATABASE <database_name> OS USER '<username>' OS GROUP '<groupname>'`

→ Tip

You can also assign OS users and groups in the *Manage Databases* app of the SAP HANA cockpit.

5. Start all tenant databases.

In the system database, execute the SQL statement `ALTER SYSTEM START DATABASE <database_name>`

→ Tip

You can also start tenant databases in the *Manage Databases* app of the SAP HANA cockpit.

Results

The system is now running in high isolation mode. As a result:

- The processes of individual tenant databases run under dedicated OS users belonging to dedicated OS groups and the processes of the system database run under the `<sid>adm` user.
- Internal database communication is authenticated using X.509 client certificates. Depending on how SSL for internal communication is configured, data communication within databases may also be encrypted. For more information about secure internal communication, see the *SAP HANA Security Guide*.
- Operations that require operating system access are restricted to users with the correct permissions. For more information, see the section on file and directory permissions with high isolation.
- New tenant databases can only be created if a dedicated OS user and group exist.

Related Information

[Database Isolation \[page 204\]](#)

[Start a Tenant Database \[page 213\]](#)

[Create a Tenant Database \[page 210\]](#)

[Assign the OS User and Group for High Isolation \[page 207\]](#)

[File and Directory Permissions with High Isolation \[page 206\]](#)

[Isolation Level High for Backups and Third-Party Backup Tools \[page 1308\]](#)

6.3.1.2.1 Database Isolation

Every tenant database is self-contained and isolated in terms of users, database catalog, repository, logs, and so on. However, to protect against unauthorized access at the operating system (OS) level, it's possible to increase isolation further through OS user separation and authenticated communication within databases.

OS User Separation

By default, all database processes run under the default OS user `<sid>adm`. If it's important to mitigate against cross-database attacks through OS mechanisms, you can configure the system for high isolation. In this way, the processes of individual tenant databases must run under dedicated OS users belonging to dedicated OS groups, instead of all database processes running under `<sid>adm`. Database-specific data on the file system is subsequently protected using standard OS file and directory permissions.

i Note

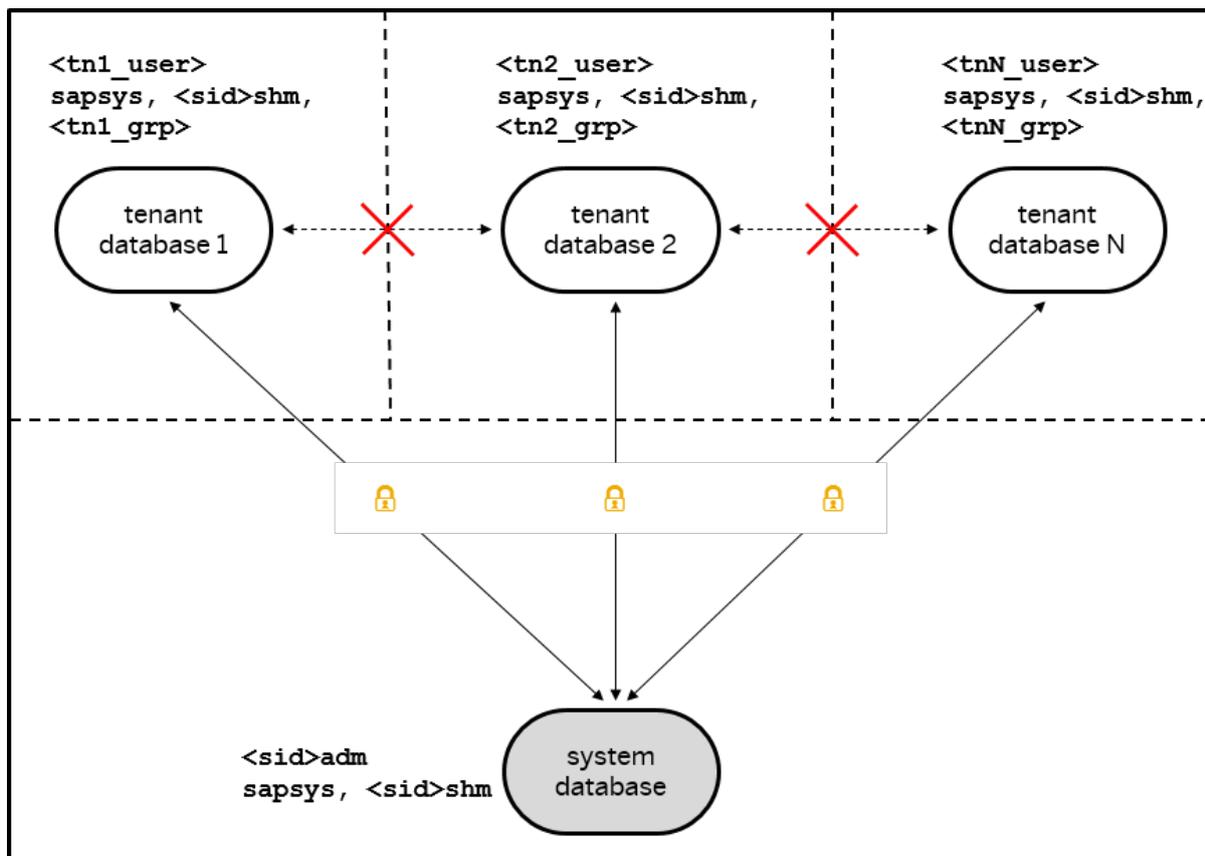
`<sid>adm` is the OS user for the system database.

Authenticated Communication

In addition, once high isolation has been configured, internal database communication is secured using the Transport Layer Security (TLS)/Secure Sockets Layer (SSL) protocol. Certificate-based authentication is used to ensure that only the processes belonging to the same database can communicate with each other. It is also possible to configure internal communication so that all data communication within databases is encrypted.

i Note

If cross-database access is enabled, communication between configured tenant databases is allowed.



High Database Isolation

Configuration

You can specify the isolation level of the system during installation. The default isolation level is low. It is also possible to change the isolation level of an existing system (from low to high or from high to low) at any time. For more information about how to do this, see *Increase the System Isolation Level* in the *SAP HANA Administration Guide*. Once high isolation has been configured, a dedicated OS user and group must exist for every tenant database. Otherwise, it's not possible to create or start a tenant database.

Internal database communication is secured with the same mechanism used for securing other internal SAP HANA communication channels. Once high isolation has been configured, authenticated communication within databases is enabled without any change required to the default TLS/SSL configuration for internal communication. However, encryption of data communication may need to be configured explicitly.

Related Information

[File and Directory Permissions with High Isolation \[page 206\]](#)

[Increase the System Isolation Level \[page 202\]](#)

[SAP HANA Administration Guide \[page 10\]](#)

6.3.1.2.2 File and Directory Permissions with High Isolation

In an SAP HANA system configured for high isolation, database-specific data on the file system is protected using standard file and directory permissions. All file and directory permissions are managed by the SAP HANA system and do not need to be set by the administrator.

System Database

The following table shows who has access to which data on the file system:

Files and Directories	Tenant OS User in Tenant OS Group	<sid>adm User
Files in directory containing system configuration files	Read permission (644)	Read permission (644)
Files in trace directory of the system database		Read and write permissions (600)
Directory containing Backint parameter file	Read permission (700)	Read permission (700)
Backint parameter file	Read and write permissions (600)	Read and write permissions (600)

Tenant Database

The following table shows who has access to which data on the file system:

Note

If you want to grant the system administrator access to the tenant database backup files and directories, you need to add the <sid>adm user to each tenant's operating system group.

Files and Directories	Tenant OS User in Tenant OS Group	<sid>adm User
Database-specific directories containing: <ul style="list-style-type: none"> Data volumes Log volumes Log mirror volumes Backups 	Read, write, and execute permissions (770)	
Database-specific directories containing: <ul style="list-style-type: none"> Configuration (*.ini) files Trace files 	Read, write, and execute permissions (770)	Read, write, and execute permissions (770)

Files and Directories	Tenant OS User in Tenant OS Group	<sid>adm User
Files in database-specific directory containing: <ul style="list-style-type: none"> • Configuration (*.ini) files • Trace files 	Read and write permissions (666)	Read and write permissions (666)
Directory containing Backint parameter file	Read, write, and execute permissions (750)	Read, write, and execute permissions (750)
Backint parameter file	Read and write permissions (640)	Read and write permissions (640)

Related Information

[Working with Third-Party Backup Tools \[page 1303\]](#)

6.3.1.2.3 Assign the OS User and Group for High Isolation

Specify the appropriate operating system user and group when moving a tenant database from a low to a high isolation level.

Prerequisites

The operating system user and operating system group exist.

Context

If you have modified the isolation level to high, during the process the entire system is restarted. Any tenant that does not have a specified OS user and OS group will not be able to restart. Perform the following in order to restart the tenant:

Procedure

1. Open *Manage Databases* in the SAP HANA cockpit by drilling down from *Overall Tenant Statuses* in the system *Overview* for the system database.
2. Select the tenant database to which you want to assign an OS user and group. When the system is running in high isolation mode, any tenant without an OS user and OS group displays a status of Not Running.

3. From the overflow menu above the table, select *Assign OS user & group*.
4. Enter the name of the existing OS user and OS group you want the tenant to use, and click *OK*.

Next Steps

Start the tenant database.

Related Information

[Clear the OS User and Group when Decreasing Isolation \[page 210\]](#)

6.3.1.3 Decrease the System Isolation Level

If you configured a system for high isolation during installation or later, you can decrease it back to the default low level if necessary. With low isolation, the processes of all databases run under the default operating system (OS) user `<sid>adm`.

Prerequisites

- You have root access to the SAP HANA system.
- You are logged on to the system database in the SAP HANA cockpit.
- You have the system privilege DATABASE ADMIN.

Procedure

1. Stop all tenant databases in the system.

In the system database, execute the SQL statement `ALTER SYSTEM STOP DATABASE <databasename>`.

→ Tip

You can also stop tenant databases in the *Manage Databases* app of the SAP HANA cockpit.

2. Configure the system for low isolation.

As the operating system user `<sid>adm`, log on to the server on which the master index server is running and run the following command:

```
python /usr/sap/<SID>/HDB<instance>/exe/python_support/convertMDC.py --
change=databaseIsolation --isolation=low
```

This command runs the following actions:

- Stops the system
 - Changes the value of the `[multidb] database_isolation` property in the `global.ini` file to `low`
 - Starts the system
3. Clear the assignment of OS users and groups to tenant databases.

In the system database, execute the SQL statement `ALTER DATABASE <database_name> OS USER '' OS GROUP ''` for every tenant database.

→ Tip

You can also clear the OS users and groups in the *Manage Databases* app of the SAP HANA cockpit.

4. Start all tenant databases.

In the system database, execute the SQL statement `ALTER SYSTEM START DATABASE <database_name>`

→ Tip

You can also start tenant databases in the *Manage Databases* app of the SAP HANA cockpit.

Results

The system is now running in low isolation mode again.

- The processes of all databases run under `<sid>adm`.
- Internal database communication is not authenticated.

Related Information

[Start a Tenant Database \[page 213\]](#)

6.3.1.3.1 Clear the OS User and Group when Decreasing Isolation

If you have modified the isolation level from high to low, you may wish to clear the previously assigned operating system user and operating system group.

Procedure

1. Open [Manage Databases](#) in the SAP HANA cockpit by drilling down from [Overall Tenant Statuses](#) in the system [Overview](#) for the system database.
2. Select the tenant database from which you want to clear the OS user and group.
3. From the overflow menu above the table, select [Assign OS user & group](#).
4. Delete the name of the existing OS user and OS group, and click [OK](#).

Next Steps

Start the tenant database.

Related Information

[Assign the OS User and Group for High Isolation \[page 207\]](#)

6.3.1.4 Create a Tenant Database

You create tenant databases after installation of a multiple-container system, after conversion from a single-container system to a single-tenant system, or anytime a new database is needed. You create tenant databases from the system database using [Manage Databases](#) in the SAP HANA cockpit.

Prerequisites

- You have the system privilege `DATABASE ADMIN`.
- If the system is configured for high isolation, the operating system (OS) user and group required for the new tenant database already exist. For more information, see [Increase the System Isolation Level](#) in the *SAP HANA Administration Guide*.

Procedure

1. Open *Manage Databases* in the SAP HANA cockpit by drilling down from *Overall Tenant Statuses* in the system *Overview* for the system database.
2. Select *Create Tenant*.
The *Create Tenant Database* page opens.
3. Enter the name of the new database and the password of the SYSTEM user.

Note

The password must initially comply with the password policy configured in the system database. Once the database is created, you can change the password policy for the tenant database if you want.

4. Optional: Specify the OS user and group of the tenant database.

If the system in which you are creating the tenant database is configured for high isolation, the processes of individual tenant databases **must** run under dedicated OS users in dedicated OS groups.

Note

If the system is configured for low isolation (default), all tenant database processes run under the default OS user `<sid>adm`.

5. Optional: Prevent the database from being started immediately after creation.

By default, the tenant database will be started immediately after creation. If you don't want this to happen, open the *Advanced Settings* section and deselect the *Start Automatically* option.

Example

You want to configure the new database before starting it to avoid having to restart.

6. Optional: Specify the host on which the database is to be created.

If the system is distributed across multiple hosts, you can specify on which host you want the master index server to start. You do this in the *Advanced Settings* section by selecting the host for the default service. If you don't select a host, load-balancing algorithms will determine optimal host placement.

7. Optional: Specify the number of the internal communication port of the master index server.

You do this in the *Advanced Settings* section by entering the port number for the default service. If you don't enter a port, it is assigned automatically based on port number availability. For more information about port number assignment, see *Connections for Tenant Databases* in the *SAP HANA Master Guide*.

8. Optional: Add any additionally required services.

- a. In the *Advanced Settings* section, choose *Add Service*.
- b. Select the service you want to add.
- c. Optional: Select the host and enter the port number of the new service.

If you don't select a host or enter a port number, they will be automatically determined.

9. Click *Create Tenant Database*. If the host does not have at least three available ports, a dialog prompts you with the option to reserve additional instance numbers.

The system starts creating the database. This may take a few moments to complete.

Technically, the creation process runs in the background as follows:

- The database is assigned a unique system local ID.
- If you did not specify host information, load-balancing algorithms determine optimal host placement.
- If you did not specify the number of the internal communication port, it is assigned automatically based on port number availability.
- The necessary data and log volumes are created on the affected hosts.
- The new database is entered in the M_DATABASES system view of the system database.
- The `daemon.ini` file is updated and the daemon process is triggered to start the indexserver service and any additionally added services on each configured host.
- The specified password is set for the user SYSTEM in the new database.

Results

The new tenant database is created and possibly started, and appears in [Manage Databases](#). It is now also in the M_DATABASES view (`SELECT * FROM "PUBLIC"."M_DATABASES"`).

Delivery units (DUs) containing automated content start to be deployed in the background. If the system is online, you can monitor the progress of deployment by executing the following statement:

```
SELECT * FROM "PUBLIC"."M_SERVICE_THREADS" WHERE THREAD_TYPE = 'ImportOrUpdate Content';
```

For more information about automated content, see *SAP HANA Content* in the *SAP HANA Security Guide*.

Next Steps

- Perform a full data backup. For more information, see *Performing Backups* in the *SAP HANA Administration Guide*.
- Adjust the value for the maximum number of asynchronous I/O requests by updating the value of the `fs.aio-max-nr` parameter in `/etc/sysctl.conf`. For more information, see *Linux Kernel Parameters* in the *SAP HANA Administration Guide*.
- Configure the new tenant database as required. For more information, see the section on managing Tenant Databases in the *SAP HANA Administration Guide*.

Related Information

[Linux Kernel Parameters \[page 944\]](#)

[Delete a Tenant Database \[page 216\]](#)

[Start a Tenant Database \[page 213\]](#)

[Monitoring Tenant Databases in SAP HANA Cockpit \[page 244\]](#)

[Increase the System Isolation Level \[page 202\]](#)

6.3.1.5 Start a Tenant Database

You start tenant databases from the system database on the [Manage Databases](#) page of the SAP HANA cockpit.

Prerequisites

- The database user with which you connect to the SAP HANA database has the privilege `DATABASE START` or `DATABASE ADMIN`.

Context

As a system administrator, you can start tenant databases either individually, or all at once by starting the whole system. For more information about how to start the whole system, see the sections on stopping and starting a resource.

Procedure

1. Open [Manage Databases](#) in the SAP HANA cockpit by drilling down from [Overall Tenant Statuses](#) in the system [Overview](#) for the system database.
2. Select the tenant database that you want to start.
3. Choose [Start Tenant](#).
The database starts. This may take a few moments.

Results

- The database is started.
- The status of database changes accordingly.

Related Information

[Start a Resource \[page 179\]](#)

[SAP HANA SQL and System Views Reference](#)

[Prevent the Start of a Tenant Database at System Startup \[page 239\]](#)

6.3.1.6 Stop a Tenant Database

You stop tenant databases from the system database on the [Manage Databases](#) page of the SAP HANA cockpit.

Prerequisites

- The database user with which you connect to the SAP HANA database has the privilege `DATABASE STOP` or `DATABASE ADMIN`.
- Consider backing the database up first.

Context

As a system administrator, you can stop tenant databases either individually, or all at once by stopping the whole system. For more information about how to stop the whole system, see the sections on stopping and starting a resource.

Procedure

1. Open [Manage Databases](#) in the SAP HANA cockpit by drilling down from [Overall Tenant Statuses](#) in the system [Overview](#) for the system database.
2. Select the tenant database that you want to stop.
3. Choose [Stop Tenant](#).
The database stops. This may take a few moments.

Results

- The database is stopped.

i Note

This is a hard stop. The database is stopped immediately even if users are connected. Open transactions are aborted and rolled back; no savepoint operation is forced. It is not possible to back up a stopped database.

- The status of database changes accordingly.

Related Information

[Stop a Resource \[page 180\]](#)

[Creating Backups \[page 1313\]](#)

[SAP HANA SQL and System Views Reference](#)

6.3.1.7 Rename a Tenant Database

You can use the cockpit to rename a tenant database.

Prerequisites

- The database user with which you connect to the SAP HANA database has the privilege `DATABASE ADMIN`.
- The database to be renamed is not running.
- The database to be renamed is not part of a copy or move operation.
- The database to be renamed does not have system replication active.
- The database to be renamed does not have the SAP HANA dynamic tiering option.

Procedure

1. Open [Manage Databases](#) in the SAP HANA cockpit by drilling down from [Overall Tenant Statuses](#) in the system [Overview](#) for the system database.
2. Select the tenant database that you want to rename.
3. If necessary, choose [Stop Tenant](#).
The database stops. This may take a few moments.
4. From the overflow menu above the table, select [Rename Tenant](#).
5. In the [Rename Tenant](#) dialog, enter the new name, and select [Rename](#).

The database is renamed. On-disk directories that contain the tenant name are also renamed. Existing backups are not renamed but backup history remains continuous.

Related Information

[Create a Tenant Database \[page 210\]](#)

[Delete a Tenant Database \[page 216\]](#)

6.3.1.8 Delete a Tenant Database

You can delete tenant databases that are no longer required. You delete tenant databases from the system database using the *Manage Databases* app of the SAP HANA cockpit.

Prerequisites

You have the system privilege `DATABASE ADMIN`.

Context

If you delete a tenant database that is running SAP HANA 2.0 SPS 01 or later, you have the option to keep the backup directories of the deleted tenant. Backups can then only be removed by deleting them from the file system. If you delete a tenant database that is running an earlier version of SAP HANA, the backup directories will be deleted automatically. It is therefore recommended that if you want to preserve these backup directories, you relocate them before deleting the database.

Deletion a tenant database (including all current backup directories) uses the `DROP DATABASE` statement in conjunction with the `DROP BACKUPS` clause. Backup directories that were previously in use, and backups that are written to third-party backup tools, are not deleted.

i Note

Once you have deleted the tenant, you can still access and consume any undeleted database backups by creating a new tenant with the same name. This will only work if the system was not configured for high isolation.

Procedure

1. Open *Manage Databases* in the SAP HANA cockpit by drilling down from *Overall Tenant Statuses* in the system *Overview* for the system database.
2. Stop the tenant database that you plan to delete by selecting it and then clicking *Stop Tenant* .
The system commences the process to stop the database. Once stopped, its status changes to *Not running*.
3. From the overflow menu above the table, choose *Delete Tenant* .
4. If this database is running SAP HANA 2.0 SPS 01 or later, choose whether to *Keep Backup Directories* or *Delete Directories* and proceed with the database deletion, or *Cancel* the database deletion. If the database is running an earlier version of SAP HANA, choose whether to *Delete Tenant* or *Cancel* the database deletion.

Results

The system commences the process to delete the database. Once deleted, the database disappears from the list. Volumes and trace files are removed.

Next Steps

If you configured the SAP Web Dispatcher to route HTTP(s) requests to the deleted database, you need to update the configuration.

Related Information

[Execute SQL Statements \[page 493\]](#)

[Execute SQL Statements in SAP HANA Studio \[page 118\]](#)

[Configure HTTP\(S\) Access to Tenant Databases via SAP HANA XS Classic \[page 1578\]](#)

[SAP HANA SQL and System Views Reference](#)

[Start a Tenant Database \[page 213\]](#)

6.3.1.9 Disable Features on a Tenant Database

To safeguard and/or customize your system, certain features of the SAP HANA database can be disabled in tenant databases. You can do this in the SAP HANA cockpit.

Prerequisites

- The system database is registered in the SAP HANA cockpit.
- You have the system privilege INIFILE ADMIN.

Context

Some features of the SAP HANA database are not required or desirable in certain environments, in particular features that provide direct access to the file system, the network, or other resources. To maximize your control over the security of your system, you can disable these features in tenant databases, for example import and export operations or the ability to back up the database.

The system view M_CUSTOMIZABLE_FUNCTIONALITIES provides information about those features that can be disabled and their status. This view exists in both the SYS schema of every database, where it contains

database-specific information, and in the SYS_DATABASES schema of the system database, where it contains information about the enablement of features in all databases.

For more information about the features that can be disabled and why, see *Restricted Features in SAP HANA Tenant Databases* in the *SAP HANA Security Guide*.

You disable features in tenant databases in the `customizable_functionalities` section of the `global.ini` file.

i Note

All features are enabled in the system database and cannot be disabled.

Procedure

1. Determine which feature(s) you want to disable by referring to the view M_CUSTOMIZABLE_FUNCTIONALITIES (SYS) of the system database.
2. On the [Overview](#) page of the system database in the SAP HANA cockpit, open [Configuration of System Properties](#) by clicking the corresponding administration link.
3. Select the configuration file `global.ini` file and the section `customizable_functionalities`.
4. Add a new parameter for the feature that you want to disable:
 - a. Specify the database on which you want to blacklist the properties.
 - b. In the [Key](#) field, enter the name of feature that you want to disable and set the value to `false`.

i Note

If you want to disable the feature on all tenant databases (including any that will be created in the future), enter `false` as the system layer value.

5. Repeat for further features not required in the tenant database(s).
6. Restart the affected tenant database(s).

Results

The feature is disabled. You can verify this in the view M_CUSTOMIZABLE_FUNCTIONALITIES (SYS_DATABASES).

Tenant database administrators can see which features are enabled in their database using the view M_CUSTOMIZABLE_FUNCTIONALITIES (SYS).

Related Information

[Start a Tenant Database \[page 213\]](#)

[System and Statistics Views \[page 388\]](#)

6.3.1.10 Enable and Configure Cross-Database Access

Read-only queries between tenant databases are supported but not enabled by default. You must first enable this feature for the system in the system database and then configure which databases may communicate with one another. You can do this in the SAP HANA cockpit.

Prerequisites

- The system database is registered in the SAP HANA cockpit.
- You have the system privilege INIFILE ADMIN.

Context

Every tenant database is self-contained with its own isolated set of database users and isolated database catalog. However, to support for example cross-application reporting, cross-database SELECT queries are possible. This means that database objects such as tables and views can be local to one database but be read by users from other databases in the same system.

So, for example, the following query would be possible:

```
SELECT *
FROM schema1.table1 AS tab1, db2.schema2.table2 as tab2
WHERE tab2.column2 = 'foobar'
```

For more information about which object types on remote databases can be accessed using this mechanism and which local object types can access remote database objects, see *Cross-Database Access*.

To allow queries between databases, you must first enable cross-database access and then specify which databases may communicate with one other. You can do this by configuring the `global.ini` configuration file in the SAP HANA cockpit.

Procedure

1. On the *Overview* page of the system database in the SAP HANA cockpit, open *Configuration of System Properties*
2. Select the configuration file `global.ini` file and the section by clicking the corresponding administration link `cross_database_access`.
3. Enable cross-database access by executing the following statement in the system database:

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') set
('cross_database_access', 'enabled')='true' WITH RECONFIGURE;
```
4. Enable communication from one tenant database to one or more other tenant databases by executing the following statement in the system database:

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') set
('cross_database_access',
'targets_for_<source_db_name>')='<target_db1>[,<target_db2>...]' WITH
RECONFIGURE;
```

❁ Example

You have two databases DB1 and DB2 and you want to be able to access DB1 from DB2. So you add the parameter `targets_for_DB2` with the value **DB1**.

i Note

Cross-database access is configured only in one direction. If in the above example you also want DB2 to be able to access DB1, you would have to add the parameter `targets_for_DB1` with the value **DB2**.

Results

Cross-database queries are now possible between the configured databases.

Next Steps

Create remote identities for those users who require cross-database access. For more information, see *Cross-Database Authorization in Tenant Databases* in the *SAP HANA Security Guide*.

In order for a user in one database to be able to run a query or create an object that references objects in another database, the user must be mapped to a sufficiently privileged user in the remote database.

Related Information

[Cross-Database Access \[page 220\]](#)

[Troubleshooting Error Situations Related to Cross-Database Access \[page 222\]](#)

6.3.1.10.1 Cross-Database Access

Read-only queries between tenant databases in the same SAP HANA system are possible. This supports cross-application reporting. Cross-database access must be explicitly enabled.

Every tenant database is self-contained with its own isolated set of database users and isolated database catalog. However, to support for example cross-application reporting, cross-database SELECT queries are possible. This means that database objects such as tables and views can be local to one database but be read by users from other databases in the same system.

The following object types on remote databases can be accessed using cross-database access:

- Schemas
- Rowstore and columnstore tables (not including virtual tables)
- SQL views (not including monitoring views)
- Graphical calculation views
 - If they only use supported object types as data sources
 - If they don't use procedure-based analytic privileges
- Synonyms

The following object types on the local tenant database can access database objects on the remote tenant database:

- SQL views
- Scripted and graphical calculation views
- Procedures
- Synonyms

The SAP HANA modeler supports modeling of graphical calculation views using tables and other graphical calculation views as data sources from different tenant databases. For more information, see *Tenant Databases Support for Modeling Graphical Calculation Views* in the *SAP HANA Modeling Guide (For SAP HANA Studio)*.

For more information about how to enable and configure cross-database access, see *Enable and Configure Cross-Database Access*.

Related Information

[Enable and Configure Cross-Database Access \[page 219\]](#)

[Troubleshooting Error Situations Related to Cross-Database Access \[page 222\]](#)

[Workload Management and Cross-Database Queries \[page 221\]](#)

6.3.1.10.2 Workload Management and Cross-Database Queries

Cross-database queries are executed on one or more databases. The workload management settings of the tenant database executing the query or part of the query are applied.

To balance and manage different types of workload in SAP HANA (OLAP, OLTP, mixed, and internal), it is possible to classify workloads based on user and application context information and apply resource limitations (for example, a statement memory limit). Workload classes allow SAP HANA to influence dynamic resource consumption at the session or statement level.

The execution of any plan operations of a cross-database query in a remote tenant database is subject to the resource limitations of the workload classes and mappings defined in the remote database. If multiple remote tenant databases are involved in query execution, then different limitations may apply to different portions of the execution plan.

i Note

For cross-database queries workload classes in the remote tenant database is the only way of applying resource limitations. The method of setting user parameter values for a user (statement memory limit, statement thread limit and priority are normally possible) is not supported for cross-database queries.

For more information about workload management using workload classes and workload mappings, see *Workload Management* in the *SAP HANA Administration Guide*.

Related Information

[Workload Management \[page 621\]](#)

6.3.1.10.3 Troubleshooting Error Situations Related to Cross-Database Access

If you are using cross-database access to query data from other tenant databases in your system, some error situations may arise.

Situation 1

You are creating views, procedures, or synonyms to access objects on other tenant databases in the same system. After dropping and re-creating an object on a remote tenant database, you can no longer access the view or procedure on the local tenant database. You get error messages such as `invalidated view` or `invalidated procedure`. You also notice that the `IS_VALID` column in the system views `VIEWS` and `PROCEDURES` do not accurately reflect the fact that the view or procedure is invalid. In addition, there are entries missing in the `OBJECT_DEPENDENCIES` system view for the affected views, procedures, or synonyms.

What's the problem?

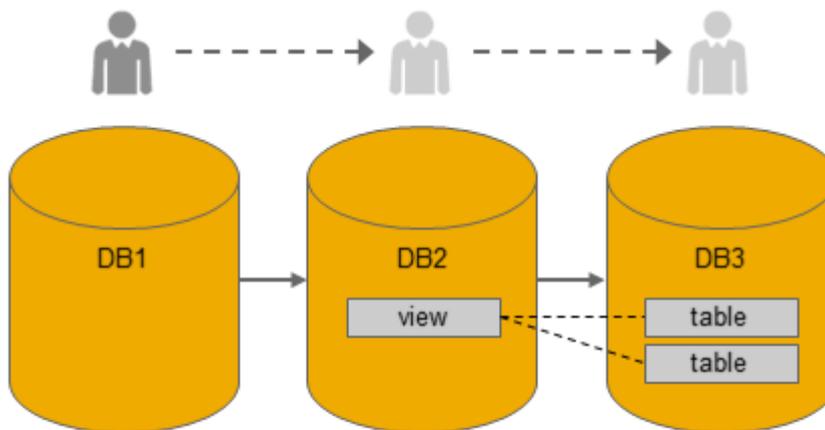
Cross-database access supports only read-only operations. Changes to an object on one tenant database cannot therefore be reflected accurately on other tenant databases that contain objects dependent on the changed object. This affects the validity flag in the relevant system views, as well as the object dependencies. Remote objects may stay valid if they retain their internal object identifier during re-creation and are re-created in an compatible way, but they will become invalid if their internal object identifier changes.

What can I do?

You need to re-create the dependent object in the local tenant database in order for it to become valid again.

Situation 2

You are querying an SQL view or a calculation view on a remote tenant database and the view itself accesses objects on a third tenant database (multi-level cross-database access). You are getting error messages such as `insufficient privilege: not authorized`. Analytic privileges on the third tenant database may be evaluated based on the wrong database user.



Executed from DB1:

 `SELECT * FROM DB2.v2`

What's the problem?

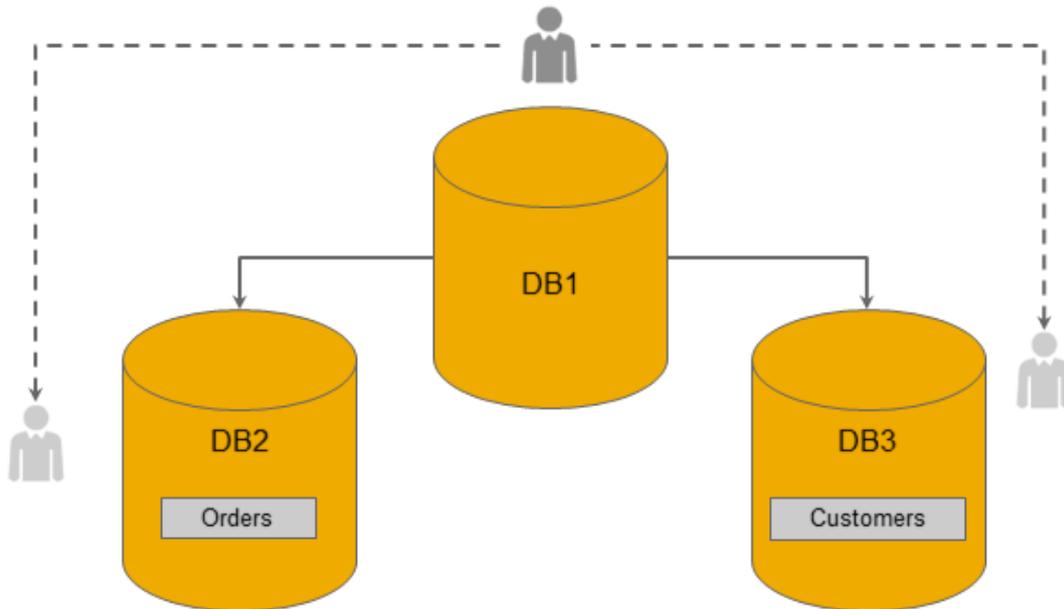
Cross-database queries do not support multiple tenant database levels as part of a view hierarchy, even if communication between databases is enabled (including the required authorized remote users).

What can I do?

Nothing. The feature is not supported.

Situation 3

Your system is running in high isolation mode. Queries that involve more than one remote tenant database run into timeouts. You are getting error messages such as `execution plan aborted` or `current operation canceled by request and transaction rolled back`. Accessing objects on remote tenant databases individually works fine.



Executed from DB1:

```

SELECT DB2.Orders.OrderID, DB3.Customers.CustomerName
FROM DB2.Orders
INNER JOIN DB3.Customers
ON DB2.Orders.CustomerID=DB3.Customers.CustomerID;

```

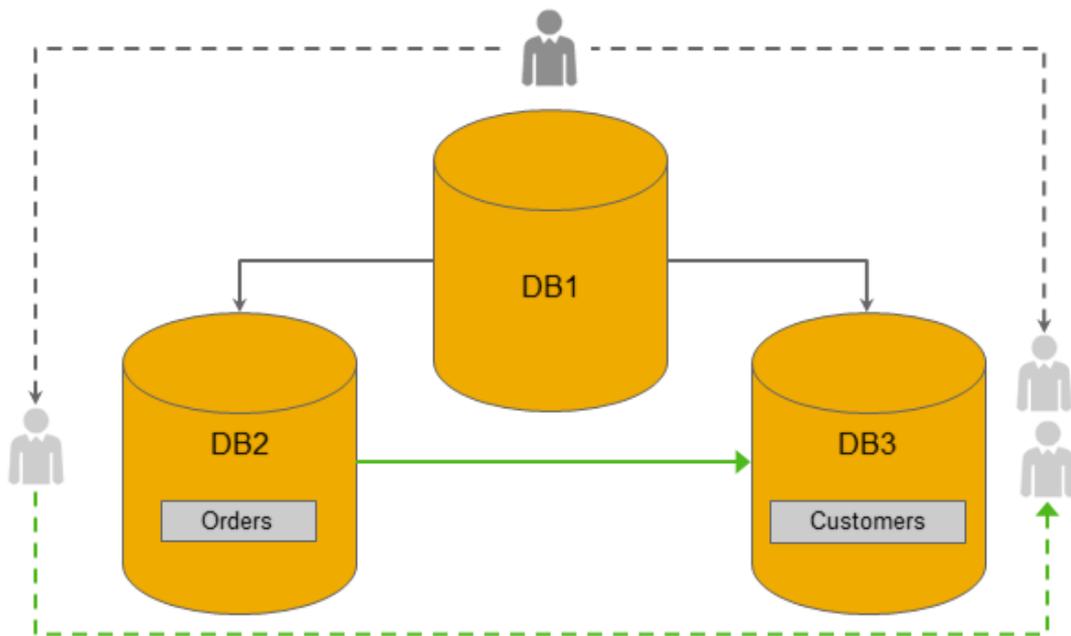


What's the problem?

The communication channels that are enabled for cross-database queries are applied in a strict fashion to the underlying network channels as well. This means that one tenant database can only open a network connection to another tenant database if communication between these two databases has been explicitly enabled. The execution plan for a query that involves objects from multiple tenant databases could however lead to direct network connections between any of the tenant databases, even if communication between them has not been explicitly enabled. This specifically applies to joins between tables on two different remote tenant databases.

What can I do?

You need to enable communication between all tenant database pairs that can potentially be involved in a query (including authorized remote users). For more information about how to do this, see *Enable and Configure Cross-Database Access*.



Executed from DB1:



```
SELECT DB2.Orders.OrderID, DB3.Customers.CustomerName
FROM DB2.Orders
INNER JOIN DB3.Customers
ON DB2.Orders.CustomerID=DB3.Customers.CustomerID;
```

Related Information

[Enable and Configure Cross-Database Access \[page 219\]](#)

6.3.1.11 Prevent Changes to System Properties in Tenant Databases

To ensure the stability and performance of the overall system or for security reasons, you can prevent certain system properties from being changed by tenant database administrators, for example, properties related to resource management. A configuration change blacklist is available for this purpose. You configure the blacklist in the SAP HANA cockpit.

Prerequisites

- The system database is registered in the SAP HANA cockpit.
- You have the system privileges INIFILE ADMIN.

Context

System configuration (*.ini) files have a database layer to facilitate the configuration of system properties for individual tenant databases. However, it may be desirable to prevent changes to certain properties being made directly in tenant databases because they could for example affect the performance of the system as a whole (CPU and memory management properties).

For this reason, a dedicated configuration change blacklist, `multidb.ini`, is available. This blacklist contains several critical properties by default. You can customize the default configuration, as well as add further properties by editing the file in the SAP HANA cockpit.

i Note

Properties in the blacklist can still be configured at all levels in the system database. For more information about configuring system properties, see *Configuring SAP HANA System Properties (INI Files)*.

Procedure

1. On the [Overview](#) page of the system database in the SAP HANA cockpit, open [Configuration of System Properties](#) by clicking the corresponding administration link.
2. Select the configuration file `multidb.ini` and the section `readonly_parameters`.
3. Add a new parameter to the blacklist:
 - a. Specify on which layer you want to blacklist the properties.

You can choose from the following layers:

Layer	Result
System	Configuration not possible in any tenant database.
Database	Configuration not possible in the specified tenant database(s)

i Note

Layered configuration is possible. A lower-layer configuration overrides a higher-layer configuration. This also allows you to change the default configuration of the blacklist. The example below shows you how you could do this.

- b. In the [Key](#) field, enter the ini file section that contains the properties you want to blacklist.

If the section exists in more than one configuration file, you can specify the exact configuration file by entering `<file>/<section>`. If you do not specify a configuration file, the properties will be blacklisted in all files that contain the section.

For example, to specify the `communication` section in all configuration files, enter `communication`. But to specify the `communication` section in the `xsengine.ini` file only, enter `xsengine.ini/communication`.

- c. In the *Value* field, enter the properties that you want to blacklist.
If you want to add all the properties in the section, enter `*`. If you want to add all the properties in all sections of a specific file, enter `<filename>/*` (for example, `xsengine.ini/*`).
- d. Choose *OK*.
- e. Add further parameters as required.

Results

Tenant database administrators cannot change the properties in the configuration change blacklist. If they try, they will get the error message: `Change not allowed for tenant database`. System administrators can still change the properties in the system database in all layers.

❖ Example

Layered Configuration

The property `[sql] sql_executors` is blacklisted for all tenant databases in all configuration files by default. You could create a layered configuration for example as follows:

- You change the `sql` entry at the system layer and enter `plan_cache_size` as the value. This overrides the default configuration so that `[sql] plan_cache_size` is blacklisted instead of `[sql] sql_executors`.
- You change the `sql` entry at the system layer and enter `sql_executors` and `plan_cache_size` as the value. This overrides the default configuration so that both `[sql] plan_cache_size` and `[sql] sql_executors` are blacklisted.
- You add a new entry `indexserver.ini/sql` at the system layer with the value `plan_cache_size` as the value. This adds a specific configuration for the `indexserver.ini` file. Here, now only `[sql] plan_cache_size` is blacklisted.

Related Information

[Configuring SAP HANA System Properties \(INI Files\) \[page 291\]](#)

6.3.1.11.1 Default Blacklisted System Properties in Tenant Databases

In systems that support tenant databases, there is configuration change blacklist `multidb.ini`, which is delivered with a default configuration.

The table below lists the system properties that are included in the `multidb.ini` file by default. This means that tenant database administrators cannot change these properties. System administrators can still change these properties in the system database in all layers.

You can customize the default configuration change blacklist by changing existing entries in the `multidb.ini` file and adding new ones. For more information about how to prevent changes to specific system properties in tenant databases in the *SAP HANA Administration Guide*.

File/Section	Properties	Description
auditing configuration	<ul style="list-style-type: none"> • <code>default_audit_trail_type</code> • <code>emergency_audit_trail_type</code> • <code>alert_audit_trail_type</code> • <code>critical_audit_trail_type</code> • <code>audit_statement_length</code> 	Prevents configuration of audit trail targets and the maximum audit statement length
communication	*	Prevents configuration of default key and trust stores, as well as other critical communication settings
global.ini/ customizable_functionalities	*	Prevents disabling of restricted features
global.ini/extended_storage	*	Prevents configuration of extended storage (SAP HANA dynamic tiering)
global.ini/persistence	<ul style="list-style-type: none"> • <code>basepath_datavolumes_es</code> • <code>basepath_logvolumes_es</code> • <code>basepath_databackup_es</code> • <code>basepath_logbackup_es</code> 	
global.ini/ system_replication	<ul style="list-style-type: none"> • <code>keep_old_style_alert</code> • <code>enable_full_sync</code> • <code>operation_mode</code> 	Prevents configuration of certain system replication settings
global.ini/ system_replication_communication	*	
global.ini/ system_replication_hostname_resolution	*	
global.ini/xb_messaging	*	Prevents configuration of messaging

File/Section	Properties	Description
multidb.ini/ readonly_parameters	*	Prevents configuration of the multidb.ini file itself
indexserver.ini/ authentication	SapLogonTicketTrustStore	Prevents configuration of the trust store for user authentication with logon/assertion tickets
memorymanager	<ul style="list-style-type: none"> • allocationlimit • minallocationlimit • global_allocation_limit • async_free_threshold • async_free_target 	Prevents configuration of memory allocation parameters
execution	max_concurrency	Prevents configuration of threading and parallelization parameters
session	<ul style="list-style-type: none"> • maximum_connections • maximum_external_connections 	
sql	sql_executors	

Related Information

[Prevent Changes to System Properties in Tenant Databases \[page 225\]](#)

[Unlock Blacklisted Parameters \[page 262\]](#)

[Copy Blacklisted Parameters \[page 263\]](#)

6.3.1.12 Configure HTTP(S) Access to Tenant Databases via SAP HANA XS Classic

To enable Web-based applications to send HTTP(S) requests to tenant databases via the SAP HANA XS classic server, the internal SAP Web Dispatcher must be configured so it knows which requests to dispatch to which database on the basis of DNS alias host names. You do this by specifying the public URL of every tenant database in the `xsengine.ini` configuration file.

Prerequisites

- You are logged on to the system database.
- You have the system privilege INIFILE ADMIN.
- The network administrator has defined an alias hostname in your organization's Domain Name System (DNS) for every tenant database in the SAP HANA system. The alias hostname must refer to the hostname of the machine that is used for HTTP(S) access to the tenant database.

- You have a role based on the role template `sap.hana.xs.wdisp.admin::WebDispatcherAdmin`. This is required to access the SAP HANA Web Dispatcher Administration tool for configuring HTTPS.

Context

The XS classic server allows Web-based applications to access SAP HANA via HTTP(S). The internal Web Dispatcher of the SAP HANA system manages these incoming HTTP(S) requests. To allow applications to send requests to specific databases, every tenant database needs an alias host name. Requests to the alias host name can then be forwarded to the XS server of the corresponding tenant database. Requests with the physical host name in the HTTP host header are forwarded to the XS server running on the system database.

The default HTTP ports are used in all cases, that is, 80<instance> (HTTP) and 43<instance> (HTTPS). Alias host names are mapped to internal HTTP(S) ports so that incoming requests can be routed to the correct database.

You configure HTTP(S) access to tenant databases by specifying in the `xsengine.ini` file the URLs by which each tenant database is publicly accessible. The system then automatically configures the Web Dispatcher by generating the required profile entries in the `webdispatcher.ini` configuration file. It is not necessary to specify the URL of the system database, this is done automatically.

i Note

This automatic configuration of the Web Dispatcher is controlled by the parameter `[profile] wdisp/system_auto_configuration` in the `webdispatcher.ini` configuration file. If this parameter is set to **false**, you need to configure the `webdispatcher.ini` file manually.

For HTTPS access, you must subsequently configure the required client certificates and trust stores using the SAP Web Dispatcher Administration tool. The following approaches are supported:

- Using a single "wildcard" server certificate in a single trust store that covers all databases in the system. Wildcard certificates are more flexible when tenant databases are frequently added and deleted. However, if you use a wildcard certificate, either the server requires its own sub-domain or you must ensure that the certificate cannot be abused from other servers.

⚠ Caution

Do not use a wildcard server certificate if strict isolation between tenant databases is required. If authentication relies on a wildcard certificate and a shared trust store, users of one tenant database will be able to log on to other databases in the system.

- Using individual certificates in individual trust stores for each database. Individual certificates for each database are more suitable in a flat domain structure for individual servers. They also ensure strict isolation between tenant databases. However, they involve more administrative effort to maintain.

Procedure

1. Specify the public URLs of all tenant databases in the `xsengine.ini` file in one of the following ways:

Option	Description
SAP HANA studio	<ol style="list-style-type: none"> 1. Open the Administration editor and choose the <i>Configuration</i> tab. 2. Navigate to the <code>xsengine.ini</code> file and expand the <code>public_urls</code> section. 3. For each tenant database in the system, add the new properties <code>http_url</code> and <code>https_url</code> at the database layer and enter its public URL as the value: <code>http://<virtual_hostname>:80<instance></code>
SQL	For each tenant database, execute the statements: <ul style="list-style-type: none"> ○ ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', '<tenant_DB_name>') SET ('public_urls', 'http_url') = 'http://<virtual_hostname>:80<instance>' WITH RECONFIGURE; ○ ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', '<tenant_DB_name>') SET ('public_urls', 'https_url') = 'https://<virtual_hostname>:43<instance>' WITH RECONFIGURE;

Note

The following values are set at the **default layer** and represent the URLs of the system database:

- `http://$ (SAPLOCALHOST) : 80$ (SAPSYSTEM)`
- `https://$ (SAPLOCALHOST) : 43$ (SAPSYSTEM)`

By default, the system database initially retrieves any request with the port `80<instance_no>`. However, as soon as you configure the URLs of tenant databases, it is available under `http://<localhost>:80<instance>` only, and not the fully qualified domain name (FQDN). The local host is known to SAP HANA without the FQDN.

If you want to change this default behavior and configure a different URL for the system database, you can do so by executing the following statement:

```
ALTER SYSTEM ALTER CONFIGURATION ('nameserver.ini', 'system')
SET('public_urls', 'http_url') = 'http://<virtual_hostname>:80<instance>'
WITH RECONFIGURE;
```

New entries are now created in the `webdispatcher.ini` file at the host layer for every database. You can verify this by executing the following statement (from the system database):

```
SELECT KEY, VALUE, LAYER_NAME FROM SYS.M_INIFILE_CONTENTS WHERE FILE_NAME =
'webdispatcher.ini' AND SECTION = 'profile' AND KEY LIKE 'wdisp/system%'
```

This returns the following result for example:

```
KEY          | VALUE                                                                 | LAYER_NAME
wdisp/system_0 | GENERATED, SID=SYS, EXTSRV=http://localhost:30014, SRCVHOST='myhost' | DEFAULT
wdisp/system_1 | GENERATED, SID=MYD, EXTSRV=http://localhost:30042, SRCVHOST='mydatabase.example.com' | HOST
```

2. Optional: Secure incoming communication by configuring HTTPS.

Option	Description
Single certificate for all databases	<ol style="list-style-type: none"> 1. Start the SAP HANA Web Dispatcher Administration tool at <code>http://<localhost>:80<instance>/sap/hana/xs/wdisp/admin/</code>. 2. For the default <code>SAPSSLS.pse</code> trust store, create a new SSL key pair and certificate request: <ol style="list-style-type: none"> 1. From the main menu, choose SSL and Trust Configuration > PSE Management. 2. From the <i>Manage PSE</i> menu, choose <i>SAPSSLS.pse</i>. 3. Choose <i>Recreate PSE</i>. 4. Enter a distinguished name that matches the host name of all tenant databases. <div data-bbox="478 577 1394 828" style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p>❖ Example</p> <ul style="list-style-type: none"> ○ Physical host name: myhost.example.com ○ Tenant host name 1: mydatabase1.example.com ○ Tenant host name 2: mydatabase2.example.com <p>In this case, you specify CN=*.example.com as the DN, thus creating a server certificate that matches all tenant databases and the system database.</p> </div> <ol style="list-style-type: none"> 5. Choose <i>Create</i>. 6. Create a certificate request and submit to your certificate authority (CA) for signing (<i>Create CA Response</i>). 3. Import the signed certificate <p>For more information, see <i>Configure HTTPS (SSL) for Client Application Access</i>.</p>
Individual certificates for each database	<ol style="list-style-type: none"> 1. Start the SAP HANA Web Dispatcher Administration tool at <code>http://<localhost>:80<instance>/sap/hana/xs/wdisp/admin/</code>. 2. For each tenant database and the system database, create a new trust store with a unique certificate: <ol style="list-style-type: none"> 1. From the main menu, choose SSL and Trust Configuration > PSE Management. 2. On the PSE management screen, choose <i>Create New PSE</i>. 3. Enter a file name for the new PSE. <div data-bbox="478 1261 1394 1366" style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p>❖ Example</p> <p><code>example.pse</code></p> </div> <ol style="list-style-type: none"> 4. Enter the distinguished name: <p>CN=<host name used for the tenant database in the public_urls section of the xsengine.ini file></p> 5. Choose <i>Create</i>. 6. For the new PSE, create a certificate request and submit to your CA for signing (<i>Create CA Response</i>). 7. Import the signed certificate into the new PSE (<i>Import CA Response</i>). 3. Configure the Web Dispatcher to use multiple certificates: <ol style="list-style-type: none"> 1. In the <code>webdispatcher.ini</code> file, create or change the parameter <code>[profile] icm/ssl_config_0</code>, specifying as the value: <p>ID=ssl_config_main, CRED=SAPSSLS.pse, SNI_CREDS=<semicolon (';') separated list of database PSE files></p> 2. Add <code>,SSLCONFIG=ssl_config_main</code> to the value of the <code>icm/server_port</code> parameter for the HTTPS port (by default <code>icm/server_port_1</code>).

Option	Description
	<p>❖ Example</p> <pre>icm/server_port_1 = PROT=HTTPS,PORT=4443\$(SAPSYSTEM),PROCTIMEOUT=600, SSLCONFIG=ssl_config_main</pre>

Results

You can access the XS server of tenant databases via the configured URLs.

→ Tip

If you experience slow response times when accessing the XS server of a tenant database (for example, Web-based applications running on the tenant database), this indicates that the server is not able to resolve the host name correctly using the DNS and retries repeatedly. If this is the case, contact your network administrator for a detailed problem analysis.

As a workaround, you can manually override virtual host name resolution on the machine where the browser is running by modifying the `/etc/hosts` file on the local machine. In this file, append a new line, starting with the static IP address of the server, followed by the virtual host name of your tenant database, for example, "10.20.30.40 mydatabase.example.com". To edit this file you need admin or root privileges.

Next Steps

Optional: Enable access to Web-based applications from the SAP HANA studio.

Some Web-based tools are accessible from the SAP HANA studio, for example, the SAP HANA cockpit and SAP HANA Lifecycle Management tool. If you want to be able to access these tools from a tenant database registered in the studio, you must specify the alias hostname in the properties. You can do this as follows:

1. In the *Systems* view, right-click the tenant database and choose *Properties*.
2. Open the *XS Properties* page and enter the alias hostname in the *XS Host* field.

Related Information

[Configure HTTPS \(SSL\) for Client Application Access \[page 1575\]](#)

[Using SAP Web Dispatcher for Load Balancing with Tenant Databases \[page 280\]](#)

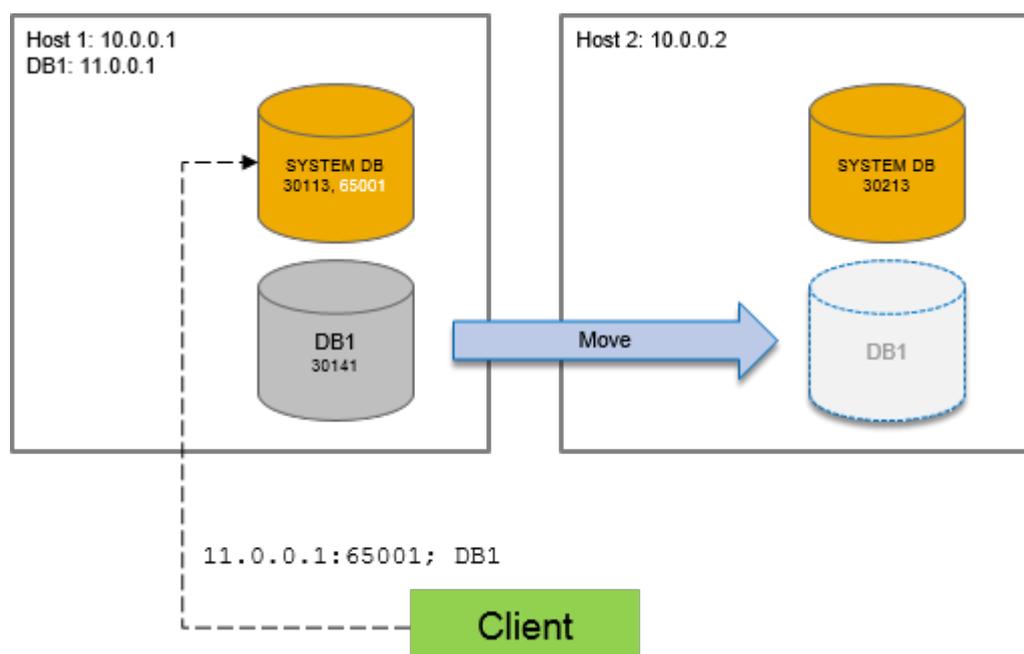
6.3.1.13 Configure Host-Independent Tenant Addresses

You can configure the access to tenant databases to be independent of the SAP HANA system ID number by mapping additional ports to a tenant database.

Context

The client connection to a tenant database is established over port 3<instance_no>13 of the system database. If a tenant database is moved to another system, the instance number of the system and consequently the connection port will change. To establish a connection independent of its current host, you can specify additional port numbers and map them to the tenants.

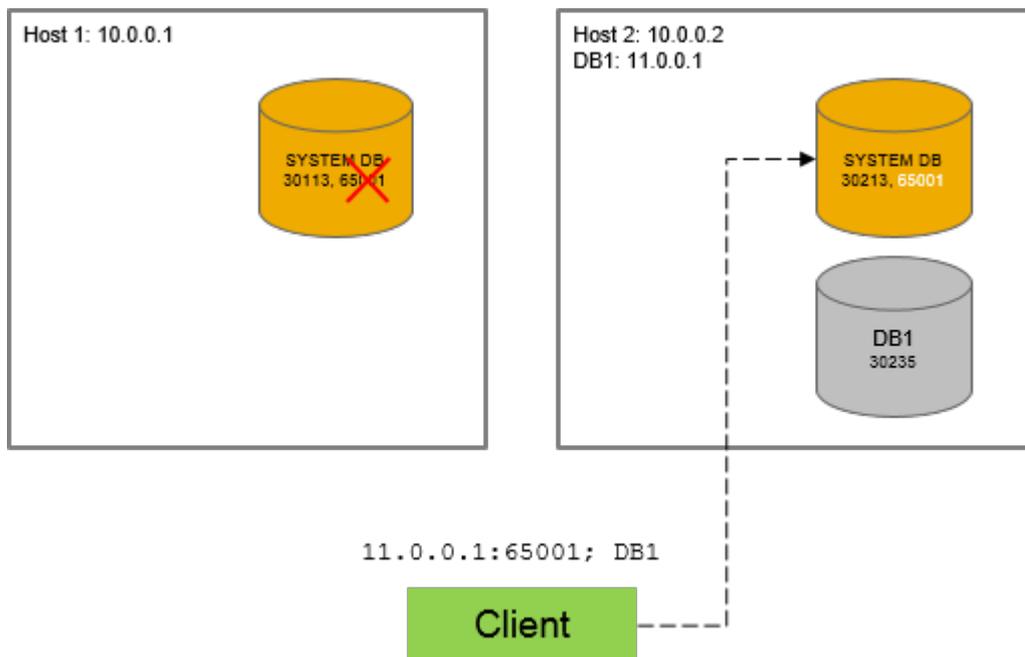
Configure a connection that is independent of the actual host IP address by mapping an IP address to each tenant database at operating system level. Add an additional port on which the system database listens for incoming connections. A client connection is then established by calling the IP address of the tenant, the name of the tenant database and the additional listen port.



Sample Code

```
SERVERNODE=11.0.0.1:65001;UID=dbuser;PWD=Aa123456;DATABASENAME=DB1
```

Once the tenant was moved to the target system, the additional listen port has to be removed on the source system and added on the target system. The tenant-specific IP address must be added for the target host at operating system level. A client connection to the tenant database can still be established with the same information as before the move.



Procedure

Configure additional ports on which the system database listens in addition to port 3<instance_no>13. Any available port number except 0 is permitted.

You can do this by executing the following SQL statement:

```
ALTER SYSTEM ALTER CONFIGURATION ('nameserver.ini' , 'system') SET ('multidb' , 'listenports' ) = '<port1>[,<port2>...]' WITH RECONFIGURE;
```

Next Steps

To remove all additional ports for a tenant database, execute the following SQL statement:

```
ALTER SYSTEM ALTER CONFIGURATION ('nameserver.ini' , 'system') UNSET ('multidb' , 'listenports' ) WITH RECONFIGURE;
```

Related Information

[Copying and Moving Tenant Databases Between Systems \[page 1004\]](#)

6.3.1.14 Create a Fallback Snapshot

You can use the cockpit to create a fallback snapshot of a tenant database. You can revert the state of a tenant database to a specific point in time if needed.

Prerequisites

- The database user with which you connect to the SAP HANA database has the privilege `DATABASE ADMIN`.

Context

You can create a fallback snapshot for a tenant database. It allows you to revert to a particular database state. If you no longer need the fallback snapshot, you can delete it.

A fallback snapshot may be useful if you perform changes to the contents of a database that you may need to roll back quickly, e.g. if you upgrade to a new version of an application.

i Note

- Fallback snapshots can only be created for tenant databases.
- Configuration changes are not included.
- You can only create one fallback snapshot per tenant database. If you need to create a new fallback snapshot, delete the existing one first.
- A service cannot be added or removed if a fallback snapshot already exists.
- A fallback snapshot cannot be created if the tenant database is the primary database in a system replication scenario.
- A fallback snapshot does not replace a database backup.
- A fallback snapshot is not included in a database backup.

Procedure

1. Open [Manage Databases](#) in the SAP HANA cockpit by drilling down from [Overall Tenant Statuses](#) in the system [Overview](#) for the system database.
2. Select the tenant database for which you want to create a fallback snapshot.
3. From the overflow menu above the table, select [Create Fallback Snapshot](#).

You can also create a fallback snapshot by executing the `ALTER DATABASE` statement in the system database:

```
ALTER DATABASE <database name> CREATE FALLBACK SNAPSHOT;
```

A fallback snapshot is created. You can verify this by querying the system view `SYS_DATABASES.M_SNAPSHOTS`.

Related Information

[Reset to a Fallback Snapshot \[page 237\]](#)

[Delete a Fallback Snapshot \[page 238\]](#)

6.3.1.15 Reset to a Fallback Snapshot

You can use the cockpit to reset a tenant database to a fallback snapshot.

Prerequisites

- The database user with which you connect to the SAP HANA database has the privilege `DATABASE ADMIN`.

Context

You can revert to a particular database state by resetting to a fallback snapshot. This may be useful if you performed changes to the contents of a database that you need to roll back quickly.

Procedure

1. Open [Manage Databases](#) in the SAP HANA cockpit by drilling down from [Overall Tenant Statuses](#) in the system [Overview](#) for the system database.
2. Select the tenant database that you want to revert to the state captured in the fallback snapshot.
3. From the overflow menu above the table, select [Reset to Fallback Snapshot](#).

You can also start a tenant database from a fallback snapshot by executing the `ALTER DATABASE` statement in the system database:

```
ALTER SYSTEM START DATABASE <database name> FROM FALLBACK SNAPSHOT;
```

The tenant database is reset to the state captured in the fallback snapshot.

Next Steps

The fallback snapshot will remain available after the reset. If you no longer need the fallback snapshot, delete it.

Related Information

[Create a Fallback Snapshot \[page 236\]](#)

[Delete a Fallback Snapshot \[page 238\]](#)

6.3.1.16 Delete a Fallback Snapshot

You can use the cockpit to delete a fallback snapshot of a tenant database.

Prerequisites

- The database user with which you connect to the SAP HANA database has the privilege `DATABASE ADMIN`.

Context

You can delete a fallback snapshot for a tenant database. You can only create one fallback snapshot per tenant database. If you need to create a new fallback snapshot, delete the existing one first.

Procedure

1. Open *Manage Databases* in the SAP HANA cockpit by drilling down from *Overall Tenant Statuses* in the system *Overview* for the system database.
2. Select the tenant database for which you want to delete a fallback snapshot.
3. From the overflow menu above the table, select *Delete Fallback Snapshot*.

You can also delete a fallback snapshot by executing the `ALTER DATABASE` statement in the system database:

```
ALTER DATABASE <database name> DROP FALLBACK SNAPSHOT;
```

The fallback snapshot is deleted.

Related Information

[Create a Fallback Snapshot \[page 236\]](#)

[Reset to a Fallback Snapshot \[page 237\]](#)

6.3.1.17 Prevent the Start of a Tenant Database at System Startup

You can prevent the start of individual tenant databases.

Prerequisites

- You are logged on to the system database.
- You have the system privilege DATABASE ADMIN.

Context

By default, all tenant databases that were running before the SAP HANA system was stopped are restarted upon system startup. For troubleshooting purposes you may want to prevent a particular database from starting until the issue is resolved. You do this in the SAP HANA cockpit using [Manage Databases](#), or from the system database using the ALTER DATABASE statement.

Procedure

You can prevent a tenant restart through the cockpit:

1. Open [Manage Databases](#) in the SAP HANA cockpit by drilling down from [Overall Tenant Statuses](#) in the system [Overview](#) for the system database.
2. Select the tenant.
3. From the overflow menu above the table, choose [Set Restart Mode](#).
4. In the dialog, select [No auto-restart](#) from the drop-down menu.

The tenant database will not be started after a system restart. You can opt to display the restart mode of all tenants through the gear icon in the [Manage Databases](#) table header.

Alternatively, you can prevent a tenant restart by executing the ALTER DATABASE statement:

```
ALTER DATABASE database_name NO RESTART
```

The tenant database will not be started after a system restart. You can verify this by querying the public view M_DATABASES. The result will look like this:

DATABASE_NAME	DESCRIPTION	RESTART_MODE
SYSTEMDB	SystemDB-<SID>-<INSTANCE>	DEFAULT
<SID>	SingleDB-<SID>-<INSTANCE>	NO

To restore the default behavior, execute the following ALTER DATABASE statement:

```
ALTER DATABASE <database_name> DEFAULT RESTART
```

Then, start the tenant database manually. At the next system startup, the tenant database will be restarted.

Related Information

[Start a Tenant Database \[page 213\]](#)

[Stop a Tenant Database \[page 214\]](#)

6.3.1.18 Copy or Move a Tenant Database Using Replication

Use system replication to copy or move a tenant database from one system to another.

Prerequisites

- For the source (the original tenant database): if encryption is required for the copy or move, you need an account with the privilege `INIFILE ADMIN`.
- For the target (the resource where you're putting the moved or copied database):
 - You need an account with the privileges `DATABASE ADMIN`, `CREDENTIAL ADMIN`, and `CATALOG READ`.
 - If encryption is required for the copy or move, you also need the privileges `CERTIFICATE ADMIN` and `TRUST ADMIN` for the target resource.
- Back up the tenant database before copying or moving it.

⚠ Caution

When you use the cockpit to move a tenant, the source database is deleted as part of the process. If the source is running SAP HANA 2.0 SP01 or earlier, its backups are also deleted as part of the process—you can't roll back! Before moving, SAP recommends that you run a backup, then replicate the backup to a new location.

Context

For conceptual background information, refer to *Copying and Moving Tenant Databases Between Systems* in the *SAP HANA Administration Guide*.

Procedure

1. On the [Overview](#) page for the system database, click [Overall Tenant Statuses](#) to drill down to the Manage Databases page.
2. In the overflow menu (top right), select [Configure Replication](#).
3. Use the drop-down menus to choose the source and target resources.
If encryption is required for the copy or move, you'll see a notice warning that both the source and the target will be restarted.
4. Enter credentials if the cockpit prompts you to do so. The cockpit alerts you about any missing privileges.
5. Enter or browse to the location of the public key certificate if the cockpit prompts you to do so.
6. Click [Review](#) to go over the information you've provided. You can use the [Edit](#) links on the review page to make changes.
7. Click [Prepare for Copy/Move](#) (bottom right) to continue.

i Note

If encryption is required, this step triggers the restarts of the source and target resources.

8. On the progress screen, wait for the cockpit to complete the steps needed to get ready for the copy or move operation. A check appears next to each step as it's completed. When all steps are complete, click the [Copy or move ...](#) link to continue.
9. Choose whether to copy or move the tenant database.
10. On the source page, you can change the source database. The cockpit prompts you if authentication is required or if the source has not been backed up.
11. On the target page, you can change the target database.
12. (Optional) Under [Advanced Settings](#), the cockpit lists the source's services. Use the services fields to create corresponding new services for the target. (Advanced settings are not present on a resource that has only one host.)
13. If the isolation level requires it, the cockpit prompts you to enter a dedicated OS user and group for the source.
14. Click [Review](#), check the information on the review page, and click [Copy Tenant Database](#) or [Move Tenant Database](#) (bottom right) to continue.

The progress screen shows how far along the copy or move process is.

If you change your mind before the process is complete, click [Cancel Copy](#) or [Cancel Move](#) (lower right). The cockpit stops and drops the target tenant database.

If you don't want to wait, click [Run in Background](#) (upper right) to go to the Manage Databases page for the source, where the tenant's status is [Copying](#) or [Moving](#) until the process is complete.

15. When the cockpit reports that the copy or move succeeded on the progress page, click [Go to Backup](#) and immediately back up the new tenant.

Next Steps

Register the new tenant with the cockpit or ask your administrator to do so.

Related Information

[Monitoring Tenant Databases in SAP HANA Cockpit \[page 244\]](#)

[Create Data Backups and Delta Backups \[page 1314\]](#)

[Set Up SAP HANA System Replication from the Primary System \[page 1099\]](#)

[Copying and Moving Tenant Databases Between Systems \[page 1004\]](#)

6.3.1.19 Reset the SYSTEM Password of a Tenant using the Cockpit

If the password of the SYSTEM user in a tenant database is lost, you as the system administrator can reset it from the system database.

Prerequisites

- You cannot log on to the tenant database as the SYSTEM because the password has been irretrievably lost.
- There is no user available with the system privilege USER ADMIN who can reset the SYSTEM user password.

i Note

If you can log on as SYSTEM or another user with the system privilege USER ADMIN, do not use the procedure described here to change the password of the SYSTEM user. Instead, change the password using the *User* editor in the SAP HANA cockpit

- You are connected to the system database and have the system privilege DATABASE ADMIN.

Procedure

1. Open *Manage Databases* in the SAP HANA cockpit by drilling down from *Overall Tenant Statuses* in the system *Overview* for the system database.
2. Stop the tenant database by selecting it and then clicking *Stop Tenant* .
The system commences the process to stop the database. Once stopped, its status changes to *Not running*.
3. From the overflow menu above the table, select *Reset Password*.
4. In the dialog, enter and confirm a new temporary password for the SYSTEM user.
5. Select *Reset Password & Restart*.

Results

- The password for the SYSTEM user is reset and the tenant database is started.
- You will have to change the password the next time you log on with this user, this time in line with the password policy of the tenant database.
- If the SYSTEM user was previously deactivated, locked, or expired, it is now activated again. In this case, we recommend that you return it to its deactivated state.
- If auditing is enabled, the password change is automatically logged in both the system and tenant database audit trails.

Related Information

[Monitoring Tenant Databases in SAP HANA Cockpit \[page 244\]](#)

[Resetting the SYSTEM User Password \[page 711\]](#)

6.3.2 Monitoring and Managing Tenant Databases

To ensure the overall performance and stability of an SAP HANA system, you as the system administrator can monitor all tenant databases in the system using the system database. You also can perform administration tasks such as stopping and starting tenant databases, or adding and removing services.

i Note

Administration of tenant databases is possible using the SAP HANA cockpit. However, command-line tools are required for some tasks.

Support Model

The following is the general approach for analyzing and resolving issues in tenant databases:

1. Tenant database administrators analyze issues in their tenant databases using the available diagnosis and trace files.
2. If tenant database administrators discover issues that they cannot analyze using diagnosis and trace files, they contact the system administrator.
3. The system administrator can first check the health of the tenant database in the system database by analyzing the monitoring data available in the SYS_DATABASES schema.
4. If the system administrator cannot see what the problem is from the system database, the tenant database administrator needs to provide him with the necessary privileges to access the tenant database directly so that the system administrator can analyze the issue there.

Related Information

[Administration of Tenant Databases \[page 20\]](#)

[Start a Tenant Database \[page 213\]](#)

[Delete a Tenant Database \[page 216\]](#)

[View Diagnosis Files of an Unavailable Tenant Database \[page 264\]](#)

[Add or Remove Services in a Tenant Database \[page 255\]](#)

6.3.2.1 Monitoring Tenant Databases in SAP HANA Cockpit

As the tenant database administrator, you can monitor the availability, resource usage, and performance of tenant databases in the SAP HANA cockpit from the system database.

Aggregate system information is available on the [Overview](#) page of the system database. The [Overview](#) page of the system database has separate sections for [Tenant Monitoring and Administration](#) and [System DB Monitoring and Administration](#). The latter displays information and links for the monitoring the system database as a resource itself. For more information, see [Using the System Overview](#).

Drilling down on the [Overall Tenant Statuses](#) tile displays the [Manage Databases](#) page which provides you with further information for all tenant databases. From [Manage Databases](#), you can then drill down to more detailed information about each individual tenant database. You can use the cockpit to monitor and manage more than one resource, each running version SAP HANA 1.0 SPS 12, or later. Any resource running version SAP HANA 2.0 SPS 01, or later is set in multiple-container mode, by default. The cockpit can also monitor single-container systems running earlier versions of SAP HANA. When you drill down to the system [Overview](#), and subsequently to [Manage Services](#), the operations you have the option to perform depend on whether you have drilled down through the system database or the tenant.

Note

To perform operations on a tenant database, you have the system privilege DATABASE ADMIN.

Related Information

[Using the Overview to Manage a Resource \[page 318\]](#)

[Database Details \[page 244\]](#)

[Monitor Alerts for a Tenant Database \[page 251\]](#)

6.3.2.1.1 Database Details

The [Manage Databases](#) page provides you with detailed information about all databases, as well as several drill-down options for more detailed information about individual databases.

The table below lists the information available for databases, as well as the available drill-down option.

Column	Description	Drill-Down Option
Database Name	Name of the tenant database	Click the database name to open the Overview, from which you can drill down to the Manage Services page Manage Services allows you to analyze the status and resource usage of the individual services of the database. For more information, see Service Details . From Manage Services , you can also stop and start services.
Status	Status of the tenant database: <ul style="list-style-type: none"> • Running • Running with Issues (where at least one service is not running, or there is at least one high alert) • Not running • Starting • Stopping 	No drill-down available
Alerts	The number of high and medium priority alerts in the database	Click the number of alerts to open them on the Alerts page Alerts allows you to view and analyze alerts occurring in the database. You can view past occurrences of alerts. For more information, see Alert Details and Alert Priorities . You can also see how alerts are configured. For more information, see Alert Checker Details and Alert Checker Statuses .
<div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px;"> <p>i Note</p> <p>Only those alerts that identify situations with a potentially system-wide impact are visible, for example, the physical memory on a host is running out. Alerts that expose data in the tenant database (for example, table names) are not visible to the system administrator in the system database.</p> </div>		
Start Time	The time of the most recent start of the database	No drill-down available
Used Memory	The used memory of the database in relation to the system	Click the used memory bar to open the Performance Monitor app The Performance Monitor app allows you to visually analyze historical performance in the database across a range of related performance indicators. For more information, see Key Performance Indicators .

Column	Description	Drill-Down Option
<i>CPU Usage</i>	The CPU usage of the database in relation to the system	Click the CPU usage bar to open the <i>Performance Monitor</i> app The <i>Performance Monitor</i> app allows you to visually analyze historical performance in the database across a range of related performance indicators. For more information, see <i>Key Performance Indicators</i> .
<i>Disk Usage</i>	The disk usage of the database in relation to the system	Click the disk usage bar to open the <i>Performance Monitor</i> app The <i>Performance Monitor</i> app allows you to visually analyze historical performance in the database across a range of related performance indicators. For more information, see <i>Key Performance Indicators</i> .

Related Information

[Service Details \[page 246\]](#)

[Alert Details \[page 252\]](#)

[Alert Priorities \[page 253\]](#)

[Alert Checker Details \[page 253\]](#)

[Alert Checker Statuses \[page 254\]](#)

[Key Performance Indicators \[page 249\]](#)

6.3.2.1.2 Service Details

Manage Services provides you with detailed information about database services for an individual resource.

i Note

Not all of the columns listed below are visible by default. You can add and remove columns in the table personalization dialog, which you open by clicking the personalization icon in the table toolbar.

The table below lists the information available for services.

Column	Description
Host	Name of the host on which the service is running

Column	Description
Status	<p>The status of the service</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> • <i>Running</i> • <i>Running with Issues</i> (where at least one service is not running, or there is at least one high alert) • <i>Starting</i> • <i>Stopping</i> • <i>Not Running</i> <p>To investigate why the service is not running, you can navigate to the crash dump file created when the service stopped.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>The crash dump file opens in the <i>Trace</i> tool of the SAP HANA Web-based Development Workbench. For this, you need the role <code>sap.hana.xs.ide.roles::TraceViewer</code> or the parent role <code>sap.hana.xs.ide.roles::Developer</code>.</p> </div>
Service	Service name, for example, indexserver, nameserver, xsengine, and so on
Role	<p>Role of the service in a failover situation</p> <p>Automatic failover takes place when the service or the host on which the service is running fails.</p> <p>The following values are possible:</p> <ul style="list-style-type: none"> • <i>Master</i> The service is the active master worker. • No entry The service is a slave worker. • <i>Standby</i> The service is in standby mode. It does not contain any data and does not receive any requests.
Port	Port that the system uses for internal communication between services
Start Time	<p>Time at which the service started</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>The time is given in the timezone of the SAP HANA server.</p> </div>
CPU	<p>Mini chart visualizing the CPU usage of the service</p> <p>Clicking the mini chart opens the <i>Performance Monitor</i> for a more detailed breakdown of CPU usage.</p>

Column	Description
Memory	<p>Mini chart visualizing the memory usage of the service</p> <ul style="list-style-type: none"> • Dark green shows the service's used memory. • Light green shows the service's peak memory. • The grey stroke represents the effective allocation limit. • The light grey background represents the physical memory. <p>Clicking the mini chart opens the Memory Analysis app for a more detailed breakdown of memory usage.</p>
Used Memory (MB)	<p>Amount of memory currently used by the service</p> <p>Clicking the mini chart opens the Memory Analysis app for a more detailed breakdown of memory usage.</p>
Peak Memory (MB)	Highest amount of memory ever used by the service
Effective Allocation Limit (MB)	Effective maximum memory pool size that is available to the process considering the current memory pool sizes of other processes
Memory Physical on Host (MB)	Total memory available on the host
All Process Memory on Host (MB)	Total used physical memory and swap memory on the host
Allocated Heap Memory (MB)	Heap part of the allocated memory pool
Allocated Shared Memory (MB)	Shared memory part of the allocated memory pool
Allocation Limit (MB)	Maximum size of allocated memory pool
CPU Process (%)	CPU usage of process
CPU Host (%)	CPU usage on host
Memory Virtual on Host (MB)	Virtual memory size on the host
Process Physical Memory (MB)	Process physical memory used
Process Virtual Memory (MB)	Process virtual memory
Shrinkable Size of Caches (MB)	Memory that can actually be freed in the event of a memory shortage
Size of Caches (MB)	Part of the allocated memory pool that can potentially be freed in the event of a memory shortage
Size of Shared Libraries (MB)	Code size, including shared libraries
Size of Thread Stacks (MB)	Size of service thread call stacks
Used Heap Memory (MB)	Process heap memory used
Used Shared Memory (MB)	Process shared memory used
SQL Port	SQL port number
Process ID	Process ID

Related Information

[Memory Usage in the SAP HANA Database \[page 470\]](#)

[Analyze Memory Statistics \[page 327\]](#)

[Monitoring and Analyzing with the Performance Monitor \[page 397\]](#)

6.3.2.1.3 Key Performance Indicators

The *Performance Monitor* allows you select a range of host-level and service-level KPIs to analyze historical performance data of the SAP HANA database.

Host KPIs

KPI	Description
CPU	CPU used by all processes related to the operating system (OS)
Database resident memory	Physical memory used by all SAP HANA database processes
Total resident memory	Physical memory used by all OS processes
Physical memory size	Total physical memory
Database used memory	Memory used by all SAP HANA database processes
Database allocation limit	Memory allocation limit for all SAP HANA database processes
Disk used	Disk space used by data, log, and trace files belonging to the SAP HANA database
Disk size	Total disk size
Network in	Bytes read from the network by all processes
Network out	Bytes written to the network by all processes
Swap in	Bytes read from swap memory by all processes
Swap out	Bytes written to swap memory by all processes

Services KPIs

KPI	Description
CPU	CPU used by the database process
System CPU	CPU used by the database process relative to the operating system
Memory used	Memory used by the database process
Memory allocation limit	Effective allocation limit of the database process

KPI	Description
Handles	Number of open handles in the index server process
Ping time	Indexserver ping time including nsWatchdog request and collection of service-specific KPIs
Swap in	Bytes read from swap by the process
Open connections	Number of open SQL connections
Open transactions	Number of open SQL transactions
Blocked transactions	Number of blocked SQL transactions
Statements	Number of finished SQL statements
Active commit ID range	Range between newest and oldest active commit ID
Pending session request count	Number of pending requests
Active versions	Number of active MVCC versions
Acquired record locks	Number of acquire record locks
Read requests	Number of read requests (selects)
Write requests	Number of write requests (insert, update, and delete)
Merge requests	Number of merge requests
Column unloads	Number of table and column unloads
Active threads	Number of active threads
Waiting threads	Number of waiting threads
Total threads	Total number of threads
Active SqlExecutors	Total number of active SqlExecutor threads
Waiting SqlExecutors	Total number of waiting SqlExecutor threads
Total SqlExecutors	Total number of SqlExecutor threads
Data write size	Bytes written to data area
Data write time	Time used for writing to data area
Log write size	Bytes written to log area
Log write time	Time used for writing to log area
Data read size	Bytes read from data area
Data read time	Time used for reading from data area
Log read size	Bytes read from log area
Log read time	Time used for reading from log area
Data backup write size	Bytes written to data backup
Data backup write time	Time used for writing to data backup
Log backup write size	Bytes written to log backup
Log backup write time	Time used for writing to log backup
Mutex Collisions	Number of collisions on mutexes

KPI	Description
Read/Write Lock Collisions	Number of collisions on read/write locks

Related Information

[Memory Usage in the SAP HANA Database \[page 470\]](#)

6.3.2.1.4 Monitor Alerts for a Tenant Database

Alert situations in tenant databases may potentially impact the health of the overall system. For this reason, you as system administrator can monitor alerts occurring in individual tenant databases. You can do this from the system database in the SAP HANA cockpit.

Prerequisites

To be able to drill down to the alert information of a tenant database, you must have registered it as a resource in the cockpit.

Procedure

1. On the *My Resources* page or any group overview page, click the name of a system database (on the *Top Resources with Alerts* tile or the *Recently Accessed* tile, for example) to open that system database's *Overview* page.
In the *Tenant Monitoring and Administration* area, you see the overall tenant statuses, as well as aggregate alert information.
2. Open the alerts of a particular database by clicking the number of alerts indicated for that database in the *Top Resources with Alerts* tile.

All high and medium priority alerts occurring in all databases in the system are displayed in list format on the left. To see more detailed information about a specific alert on the right, simply select it.

→ Tip

You can also access database-specific alerts from *Alerts* in the *Manage Databases* app.

Next Steps

It may be helpful to see how alerts are configured in individual tenant databases. To navigate to the configuration of alert checkers from the *Alerts* app, click [View Alert Configuration](#) in the footer toolbar.

Related Information

[Using the Overview to Manage a Resource \[page 318\]](#)

[Manage Services \[page 321\]](#)

[Work with Alerts \[page 342\]](#)

6.3.2.1.4.1 Alert Details

When you select an alert, detailed information about the alert is displayed on the right.

The following detailed information about an alert is available:

Detail	Description
Category	The category of the alert checker that issued the alert Alert checkers are grouped into categories, for example, those related to memory usage those related to transaction management and so on.
Next Scheduled Run	When the related alert checker is next scheduled to run If the alert checker has been switched off (alert checker status <i>Switched Off</i>) or it failed the last time it ran (alert checker status <i>Failed</i>), this field is empty because the alert checker is no longer scheduled.
Interval	The frequency with which the related alert checker runs If the alert checker has been switched off (alert checker status <i>Switched Off</i>) or it failed the last time it ran (alert checker status <i>Failed</i>), this field is empty because the alert checker is no longer scheduled.
Alerting Host & Port	Name and port of the host that issued the alert In a system replication scenario, alerts issued by secondary system hosts can be identified here. This allows you to ensure availability of secondary systems by addressing issues before an actual failover. For more information about monitoring secondary systems in SAP HANA, see <i>Monitoring Secondary Sites</i> in the <i>SAP HANA Administration Guide</i> .
Alert Checker	Name and description of the related alert checker
Proposed Solution	Possible ways of resolving the problem identified in the alert, with a link to the supporting app, if available

Detail	Description
Past Occurrences of Alert	Configurable graphical display indicating how often the alert occurred in the past

Related Information

[Monitoring Secondary Sites \[page 1194\]](#)

6.3.2.1.4.2 Alert Priorities

The priority of an alert indicates the severity of the problem and how quickly action needs to be taken.

Priority	Description
Information	Action recommended to improve system performance or stability
Low	Medium-term action required to mitigate the risk of downtime
Medium	Short-term action required (few hours, days) to mitigate the risk of downtime
High	Immediate action required to mitigate the risk of downtime, data loss, or data corruption

6.3.2.1.4.3 Alert Checker Details

When you select an alert checker *Alert Configuration*, detailed information about the alert checker and its configuration is displayed on the right.

The following detailed information about an alert checker is available:

Detail	Description
Header information	The name of the alert checker, its status, and the last time it ran
Description	Description of what the alert checker does, for example what performance indicator it measures or what setting it verifies
Alert Checker ID	The unique ID of the alert checker
Category	The category of the alert checker Alert checkers are grouped into categories, for example those related to memory usage, those related to transaction management, and so on.

Detail	Description
Threshold Values for Prioritized Alerting	<p>The values that trigger high, medium, low, and information alerts issued by the alert checker</p> <p>The threshold values and the unit depend on what the alert checker does. For example, alert checker 2 measures what percentage of disk space is currently used so its thresholds are percentage values.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>Thresholds can be configured for any alert checker that measures variable values that should stay within certain ranges, for example, the percentage of physical memory used, or the age of the most recent data backup. Many alert checkers verify only whether a certain situation exists or not. Threshold values cannot be configured for these alert checkers. For example, alert checker 4 detects services restarts. If a service was restarted, an alert is issued.</p> </div>
Interval	The frequency with which the alert checker runs
Schedule Active	Indicator of whether the alert checker is running automatically at the configured interval
Proposed Solution	Possible ways of resolving the problem identified by the alert checker

Related Information

[Alert Checker Statuses \[page 254\]](#)

[Configure Alerting Thresholds \[page 348\]](#)

6.3.2.1.4.4 Alert Checker Statuses

The status of an alert checker indicates whether it is running on schedule, it has failed and been disabled by the system, or you switched it off.

Status	Description
Active	The alert checker is running on schedule.

Status	Description
Failed	<p>The alert checker failed the last time it ran (for example due to a shortage of system resources), so the system disabled it.</p> <p>The alert checker remains disabled for a specific length of time before it is automatically re-enabled. This length of time is calculated based on the values in the following columns of the table STATISTICS_SCHEDULE (_SYS_STATISTICS):</p> <ul style="list-style-type: none"> • INTERVALLENGTH • SKIP_INTERVAL_ON_DISABLE <p>Once INTERVALLENGTH x SKIP_INTERVAL_ON_DISABLE has elapsed, the alert checker is re-enabled. The default values for all alert checkers are such that failed checkers remain disabled for 1 hour. The system determines the status of every alert checker and/or whether the time to re-enablement has elapsed every 60 seconds.</p> <p>You can also re-enable the alert checker manually by switching it back on in Alert Configuration.</p>
Switched Off	<p>You switched off the alert checker schedule.</p> <p>If you want the alert checker to run again automatically, you must manually switch it back on.</p>

Related Information

[Switch Alerting Off/On \[page 350\]](#)

[Configure Alerting Thresholds \[page 348\]](#)

6.3.2.2 Add or Remove Services in a Tenant Database

You can add or remove services in the SAP HANA cockpit.

Prerequisites

The database user (with which you have connected to the resource) must have the system privilege DATABASE ADMIN.

Procedure

1. Open [Manage Databases](#) in the SAP HANA cockpit by drilling down from [Overall Tenant Statuses](#) in the system [Overview](#) for the system database.

2. Click the status of the tenant to open [Manage Services](#).
3. To add a service, select [Manage Services](#) > [Add Service](#).

You can also add a service by executing the `ALTER DATABASE` statement in the system database:

What	How
A new index server with automatic host placement and automatic port assignment	<code>ALTER DATABASE DB0 ADD 'indexserver'</code>
A new index server on a specific host accessible via a specific SQL port (30148)	<code>ALTER DATABASE DB0 ADD 'indexserver' AT LOCATION 'HOST_A:30148'</code>
A new index server on a specific host and a new XS server on an automatically selected host	<code>ALTER DATABASE DB0 ADD 'indexserver' AT 'HOST_B' ADD 'xsengine'</code>

Example

Note

The SQL port number is the internal communication port number plus 1.

- a. Select a service from the list.

The main services that you may need to add are:

- indexserver
- xsengine

Note

After database creation, the xsengine service automatically runs embedded in the (master) index server. If you add a separate xsengine service, the embedded service is stopped and removed.

- docstore
- scriptserver
- dpserver

Note

You cannot add a statisticsserver service. This always runs embedded in the master index server of a tenant database.

There may be other services available to add or remove, if you have installed an SAP HANA capability (such as SAP dynamic tiering), and the host has been added to the SAP HANA SYSTEM DB.

- b. Select a host and enter its port, or allow a host to be auto-assigned.

Note

The default port number range for tenant databases is `3<instance>40–3<instance>99`. This means that the maximum number of tenant databases that can be created per instance is 20. However, you can increase this by reserving the port numbers of further instances. In the cockpit, a

dialog will prompt you to do this, or you can configure the property `[multidb] reserved_instance_numbers` in the `global.ini` file. The default value of this property is 0. If you change the value to 1, the port numbers of one further instance are available (for example, 30040—30199 if the first instance is 00). If you change it to 2, the port numbers of two further instances are available (for example, 30040—30299 if the first instance is 00). And so on.

- New data and log volumes are created on the host and the information is entered in the system landscape information of system database.
- The service is added to the `M_SERVICES` system view.
- The service is started.

4. To remove the service, select **▶ Manage Services ▶ Remove Service ▶**.

You can also remove a service by executing the `ALTER DATABASE` statement in the system database:

❖ Example

```
ALTER DATABASE DB0 REMOVE 'indexserver' AT LOCATION 'HOST_A:30146'
```

i Note

Not all services can be removed. You cannot remove a global service, the master index server, or the primary index server on a host.

- The service is stopped and removed from the system landscape information of system database.
- The service is removed from the `M_SERVICES` system view.
- Data volumes and traces files are removed.

6.3.2.3 Change the Port of a Service in a Tenant Database

You can change the port of a service in the SAP HANA cockpit.

Prerequisites

The database user (with which you have connected to the resource) must have the system privilege `DATABASE ADMIN`.

Procedure

1. Open *Manage Databases* in the SAP HANA cockpit by drilling down from *Overall Tenant Statuses* in the system *Overview* for the system database.
2. To change the port of a service, select **▶ Manage Services ▶ Change Port ▶**.

3. In the *Change Port* dialog, select the *New Port* for the service.
4. Optional: If reserved instances are configured for the host, you can choose to select ports available on another instance.

i Note

The default port number range for tenant databases is $3<instance>40-3<instance>99$. This means that the maximum number of tenant databases that can be created per instance is 20. However, you can increase this by reserving the port numbers of further instances.

5. Optional: Start the service after changing the port.

By default, the service is not started after the port is changed. If you want the service to be started after the port is changed, enable the *Start Automatically* option.

6.3.2.4 Restrict Features Available to a Tenant Database

To safeguard and/or customize your system, certain features of the SAP HANA database can be disabled in tenant databases. You can do this in the SAP HANA cockpit.

Prerequisites

- The system database is registered in the SAP HANA cockpit.
- You have the system privilege INIFILE ADMIN.

Context

Some features of the SAP HANA database are not required or desirable in certain environments, in particular features that provide direct access to the file system, the network, or other resources. To maximize your control over the security of your system, you can disable these features in tenant databases, for example import and export operations or the ability to back up the database.

The system view `M_CUSTOMIZABLE_FUNCTIONALITIES` provides information about those features that can be disabled and their status. This view exists in both the `SYS` schema of every database, where it contains database-specific information, and in the `SYS_DATABASES` schema of the system database, where it contains information about the enablement of features in all databases.

For more information about the features that can be disabled and why, see *Restricted Features in Tenant Databases* in the *SAP HANA Tenant Databases*.

You can disable features in tenant databases in the `customizable_functionalities` section of the `global.ini` file, as well as in the SAP HANA cockpit as described here.

i Note

All features are enabled in the system database and cannot be disabled.

Procedure

1. Open *Manage Databases* in the SAP HANA cockpit by drilling down from *Overall Tenant Statuses* in the system *Overview* for the system database.
2. On the *Manage Databases* screen, click *Manage Restricted Features* (upper right).
To see the *Manage Restricted Features* button, you might need to widen the window or click the three dots in the upper right corner of the screen.
3. On the *Restricted Features for Tenants* page, the active tenant is highlighted in the left pane. Click the name of another tenant to manage its features.
4. To restrict a feature, click its checkbox. Click again to clear the selection.
5. Click *Save* to disable the features you selected.

Next Steps

Stop and restart the affected tenant database.

Related Information

[Stop a Tenant Database \[page 214\]](#)

[Start a Tenant Database \[page 213\]](#)

[Lock Parameters Against Editing for a Tenant Database \[page 259\]](#)

6.3.2.5 Lock Parameters Against Editing for a Tenant Database

To ensure the stability and performance of the overall system or for security reasons, you can prevent certain system parameters from being changed by tenant database administrators, for example, parameters related to resource management. A configuration change blacklist is available for this purpose. You configure the blacklist in the SAP HANA cockpit.

Prerequisites

- The system database is registered in the SAP HANA cockpit.
- You have the system privilege INIFILE ADMIN.

Context

System configuration (*.ini) files have a database layer to facilitate the configuration of system parameters for individual tenant databases. However, it may be desirable to prevent changes to certain parameters being made directly in tenant databases because they could, for example, affect the performance of the system as a whole (CPU and memory management parameters).

You can use the cockpit to blacklist parameters for a particular database—that is, to lock them against editing. (The blacklist is stored in `multidb.ini`.) Several parameters are blacklisted by default, so you'll see them when you visit the [Blacklisted Parameters for Tenants](#) page for a tenant. You can remove default properties from that page—that is, make them editable by the tenant—and you can add parameters—lock them so they cannot be edited.

i Note

Properties in the blacklist can still be configured at all levels in the system database. For more information about configuring system properties, see [Configuring SAP HANA System Properties \(INI Files\)](#).

Procedure

1. Open [Manage Databases](#) in the SAP HANA cockpit by drilling down from [Overall Tenant Statuses](#) in the system [Overview](#) for the system database.
2. On the [Manage Databases](#) screen, click [Manage Blacklisted Parameters](#) (upper right).
3. On the [Blacklisted Parameters for Tenants](#) page, the active tenant is highlighted in the left pane. Click the name of another tenant to manage its parameters.
4. To add a parameter to the blacklist, click [Add Parameter](#) at the top of the screen.
5. In the [Add Parameter to Blacklist](#) dialog, select the configuration file in which the parameter you want to add appears. To add a parameter that appears in more than one file, select [Any with the specified section](#).
6. Enter or select your parameter's section in the configuration file. If you start to type the section name, the cockpit offers section names that match what you've entered—click to select one. If you prefer to select the section, click the double box at the end of the [Section](#) field to see a list. You can even combine the two methods: enter a few characters to narrow the number of choices offered when you click the double box.
7. Enter or select the parameter or parameters you want to add to the blacklist. If you start to type the parameter name, the cockpit offers section names that match what you've entered—click to select one. If you prefer to select the section, click the double box at the end of the [Parameters](#) field to see a list. You can even combine the two methods: enter a few characters to narrow the number of choices offered when you click the double box.

You can enter multiple parameters if they're all in the section you specified. You can also specify new parameters.

8. Click [OK](#) to save the new blacklist entries.

Related Information

[Configuring SAP HANA System Properties \(INI Files\) \[page 291\]](#)

[Restrict Features Available to a Tenant Database \[page 258\]](#)

6.3.2.5.1 Default Blacklisted System Properties in Tenant Databases

In systems that support tenant databases, there is configuration change blacklist `multidb.ini`, which is delivered with a default configuration.

The table below lists the system properties that are included in the `multidb.ini` file by default. This means that tenant database administrators cannot change these properties. System administrators can still change these properties in the system database in all layers.

You can customize the default configuration change blacklist by changing existing entries in the `multidb.ini` file and adding new ones. For more information about how to prevent changes to specific system properties in tenant databases in the *SAP HANA Administration Guide*.

File/Section	Properties	Description
auditing configuration	<ul style="list-style-type: none">• <code>default_audit_trail_type</code>• <code>emergency_audit_trail_type</code>• <code>alert_audit_trail_type</code>• <code>critical_audit_trail_type</code>• <code>audit_statement_length</code>	Prevents configuration of audit trail targets and the maximum audit statement length
communication	*	Prevents configuration of default key and trust stores, as well as other critical communication settings
global.ini/ customizable_functionalities	*	Prevents disabling of restricted features
global.ini/extended_storage	*	Prevents configuration of extended storage (SAP HANA dynamic tiering)
global.ini/persistence	<ul style="list-style-type: none">• <code>basepath_datavolumes_es</code>• <code>basepath_logvolumes_es</code>• <code>basepath_databackup_es</code>• <code>basepath_logbackup_es</code>	
global.ini/ system_replication	<ul style="list-style-type: none">• <code>keep_old_style_alert</code>• <code>enable_full_sync</code>• <code>operation_mode</code>	Prevents configuration of certain system replication settings

File/Section	Properties	Description
global.ini/ system_replication_communi- cation	*	
global.ini/ system_replication_hostnam- e_resolution	*	
global.ini/xb_messaging	*	Prevents configuration of messaging
multidb.ini/ readonly_parameters	*	Prevents configuration of the multidb.ini file itself
indexserver.ini/ authentication	SapLogonTicketTrustStore	Prevents configuration of the trust store for user authentication with logon/ assertion tickets
memorymanager	<ul style="list-style-type: none"> • allocationlimit • minallocationlimit • global_allocation_limit • async_free_threshold • async_free_target 	Prevents configuration of memory allo- cation parameters
execution	max_concurrency	Prevents configuration of threading and parallelization parameters
session	<ul style="list-style-type: none"> • maximum_connections • maximum_external_con- nections 	
sql	sql_executors	

Related Information

[Prevent Changes to System Properties in Tenant Databases \[page 225\]](#)

[Unlock Blacklisted Parameters \[page 262\]](#)

[Copy Blacklisted Parameters \[page 263\]](#)

6.3.2.5.2 Unlock Blacklisted Parameters

Remove a parameter from the blacklist for a tenant database so that the parameter can be edited.

Prerequisites

- The system database is registered in the SAP HANA cockpit.

- You have the system privilege INIFILE ADMIN.

Context

By removing a parameter from the blacklist, you can enable it to be edited. However, you can remove only parameters that you or other users have added to the blacklist—you can't remove parameters that are on the blacklist by default. Default parameters are displayed without delete or edit controls on the [Blacklisted Parameters for Tenants](#) page.

Procedure

1. Open [Manage Databases](#) in the SAP HANA cockpit by drilling down from [Overall Tenant Statuses](#) in the system [Overview](#) for the system database.
2. On the [Manage Databases](#) screen, click [Manage Blacklisted Parameters](#) (upper right).
To see the [Manage Blacklisted Features](#) button, you might need to widen the window or click the three dots in the upper right corner of the screen.
3. On the [Blacklisted Parameters for Tenants](#) page, the active tenant is highlighted in the left pane. Click the name of another tenant to manage its parameters.
4. To remove a parameter to the blacklist, click the red X to the right of the parameter to be deleted and confirm the deletion.

The cockpit displays the blacklist without the parameter you removed.

Related Information

[Lock Parameters Against Editing for a Tenant Database \[page 259\]](#)

6.3.2.5.3 Copy Blacklisted Parameters

Copy parameters that have been added to another tenant's blacklist.

Prerequisites

- The system database is registered in the SAP HANA cockpit.
- You have the system privilege INIFILE ADMIN.

Context

Copy the blacklist from one tenant to another. The tenant receiving the copy is the one selected in the Tenants list in the left pane. If you've added parameters to the target tenant before copying, you can choose whether to keep them. Parameters on the blacklist by default are not copied.

Procedure

1. Open [Manage Databases](#) in the SAP HANA cockpit by drilling down from [Overall Tenant Statuses](#) in the system [Overview](#) for the system database.
2. On the [Manage Databases](#) screen, click [Manage Blacklisted Parameters](#) (upper right).
3. On the [Blacklisted Parameters for Tenants](#) page, the active tenant is highlighted in the left pane. Click the name of another tenant to manage its parameters.
4. Click [Copy Parameters](#) at the top of the screen.
5. In the [Copy Parameters](#) dialog, select the tenant whose blacklist you want to copy (the source), and specify the tenant to which you want to copy the parameters (the target).
6. To make this tenant's blacklist an exact copy of the source tenant's blacklist, discarding any parameters you've added, select [Replace all parameters, making target and source identical](#). To keep parameters you've added, select [Augment the target parameters with those from the source](#).
7. Click [OK](#) to copy the blacklist.

6.3.2.6 View Diagnosis Files of an Unavailable Tenant Database

If a tenant database is unavailable, for example because it's stopped or experiencing major performance problems, the tenant database administrator can't access diagnosis files. In this case, you as the system administrator can access the diagnosis files of the tenant database from the system database using the SAP HANA database explorer.

Prerequisites

Because the database explorer is integrated with the SAP Web IDE for SAP HANA and the SAP HANA cockpit, you must be a user of one of these applications.

Procedure

1. Open the database explorer from the SAP HANA cockpit.

2. Add a database by clicking the [Add Database](#) icon (+) at the top of the database browser pane on the left.
3. Open the [Host Diagnostic Files](#) folder of your cockpit resource, then click the diagnostic file that you want to examine to open it in an editor. The [Host Diagnostic Files](#) folder contains all diagnostic files that have been configured for the SAP Host Agent.

For more information about configuring the SAP Host Agent, see the *SAP Host Agent* documentation.

The cockpit resource must have valid SAP Control Credentials set in the cockpit. If the user has not set valid SAP Control Credentials, then an error is returned.

The diagnosis files of the system database are displayed.

Next Steps

If more detailed diagnosis information is required (for example for SAP Support), you can trigger the collection of a full system information dump for tenant databases. For more information, see *Collecting Diagnosis Information for SAP Support* in the *SAP HANA Administration Guide*.

Related Information

[Add SAP HANA Cockpit Resources and Databases to the SAP HANA Database Explorer \[page 53\]](#)

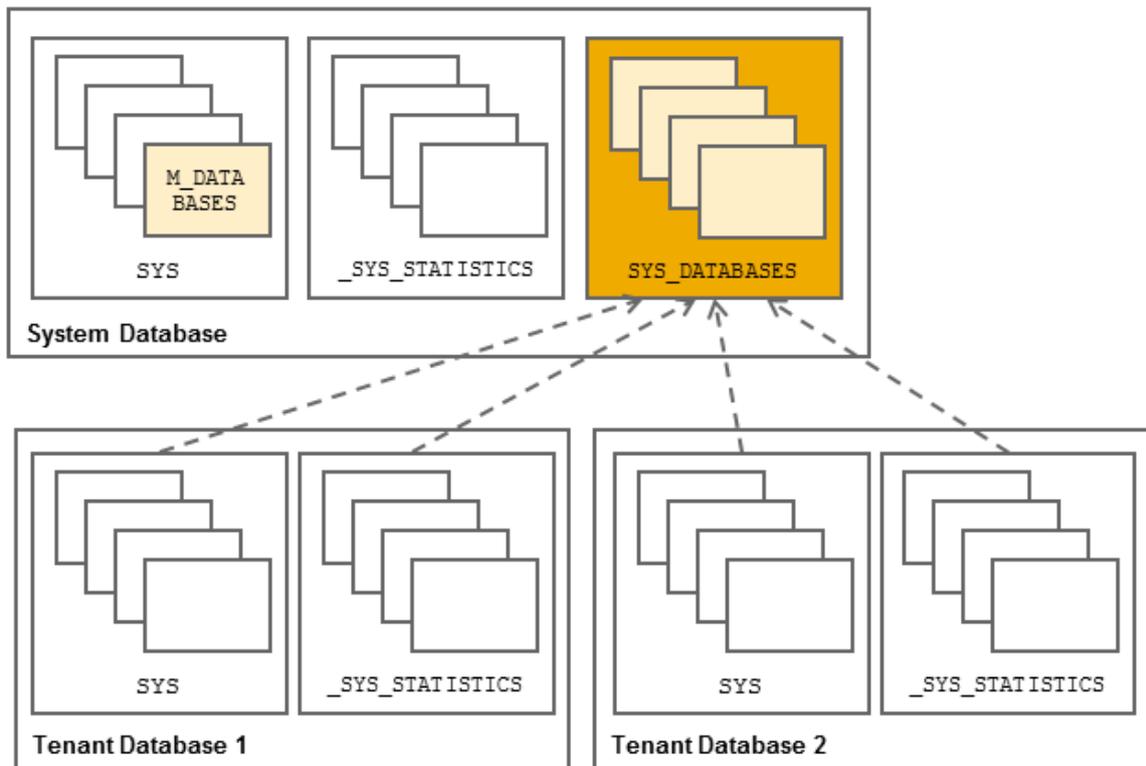
[View Diagnostic Files in the SAP HANA Database Explorer \[page 662\]](#)

6.3.2.7 System and Statistics Views in Tenant Database Systems

Every database has its own SYS and _SYS_STATISTICS schemas that contain information about that database only. For system-level monitoring, additional views are accessible in the system database: the M_DATABASES (SYS) view and the views in the SYS_DATABASES schema.

- M_DATABASES
This view is available in the SYS schema of the system database of a multiple-container system. It provides an overview of all tenant databases in the system. Only users with the system privilege DATABASE ADMIN can see the contents of this view.
- SYS_DATABASES schema
The views in the SYS_DATABASES schema provide aggregated information from a **subset** of the views available in the SYS and _SYS_STATISTICS schemas of all tenant databases in the system. These union views have the additional column DATABASE_NAME to make it possible to identify from which database the information is coming refers. The system views in the SYS_DATABASES schema are accessible only from the system database. To be able to view information in these views, you need the system privilege DATABASE ADMIN or CATALOG READ.

Tools such as the SAP HANA cockpit use these views to support system-level monitoring.



System and Statistics Views

6.3.3 Configuring Memory and CPU Usage for Tenant Databases

Manage and control the memory and CPU usage of your system by configuring limits for individual tenant databases. If necessary, you can also reserve memory for the system database.

Managing Resource Usage of Tenant Databases

Several system properties allow you to influence the allocation of memory and CPU resources in SAP HANA systems. System properties (INI) files have a database layer to facilitate the configuration of properties for individual tenant databases.

The properties listed below are particularly useful for influencing the resource consumption of tenant databases.

- `[memorymanager] allocationlimit` in the `global.ini` file
Use this property to limit the maximum amount of memory (in MB) that can be allocated individually to processes of a tenant database. Each process of a tenant database can allocate the specified value. Setting

the allocation limit too low might cause the tenant database to become inaccessible until more memory can be allocated.

❖ Example

```
Executed from the system database:ALTER SYSTEM ALTER CONFIGURATION ('global.ini',  
'DATABASE', 'MYDB') SET ('memorymanager', 'allocationlimit') = '8192' WITH  
RECONFIGURE;
```

i Note

Memory alignment will happen on the fly and may therefore take some time. To make it happen immediately, you can restart the database.

- [execution] `max_concurrency` in the `global.ini` file
Use this property to influence the maximum number of CPU cores that can be used for each tenant database by limiting the number of concurrently running threads used by the JobExecutor subsystem. A reasonable default value is the number of cores divided by the number of tenant databases. Do not specify a value of 0. A change of this value takes effect immediately.

❖ Example

```
Executed from the system database:ALTER SYSTEM ALTER CONFIGURATION ('global.ini',  
'DATABASE', 'MYDB') SET ('execution', 'max_concurrency') = '4' WITH  
RECONFIGURE;
```

i Note

In NUMA architectures, setting the `max_concurrency` parameter is not enough to achieve the desired performance gains, so you should also bind sockets that share memory using the affinity setting. For more information, see *Controlling CPU Consumption*.

Managing Memory Usage of System Database

After installation, the system database contains only data required to monitor and manage the system, as well as statistics data related to itself. This results in an average memory consumption of 15 GB.

However, if the system database is experiencing performance problems, for example, out-of-memory situations, you can reserve a minimum amount of memory (MB) for the system database by configuring the parameter `[multidb] systemdb_reserved_memory` in the `global.ini` file.

Related Information

[Controlling Parallel Execution of SQL Statements \[page 630\]](#)

[Controlling CPU Consumption \[page 277\]](#)

[Configuring SAP HANA System Properties \(INI Files\) \[page 291\]](#)

6.3.3.1 Define Memory Allocation Limits

As part of the provisioning process, you can ensure that memory is shared appropriately between tenant databases. By setting the memory allocation limits for each database on a host, you ensure appropriate memory sharing through setting the maximum amount of memory that can be allocated for a particular tenant database.

Context

In the SAP HANA cockpit you can use the *Memory Allocation* tab of the *Configure Workload Allocation* app to view and modify memory allocation limits. These limits correspond to the settings of the `allocationlimit` parameter in the memory manager section of the `global.ini` file.

SAP HANA preallocates and manages its own memory pool, used for storing in-memory table data, thread stacks, temporary results, and other system data structures. When more memory is required for table growth or temporary computations, the SAP HANA memory manager obtains it from the pool. When the pool cannot satisfy the request, the memory manager increases the pool size by requesting more memory from the operating system, up to a predefined allocation limit.

You can adjust the allocation limit for the services in the system at the following levels:

- System
- Tenant Database (these settings override the system settings)
- Host (these settings override the database and system settings)

For each database (or each database per host in a multi-host system), you can also refer to a mini-chart representing the memory usage of the `indexserver` in one day, where the vertical line shows the allocation limit setting for this specific database, dark green shows the actual used memory and light green shows the peak used memory.

The allocation limit value is applied to each service (e.g. `indexserver`, `nameserver`, `compileserver`,). However, when the *Configure Workload Allocation* app calculates the total memory available for allocation, only the `indexserver` is considered, since it requires much more memory than the other services.

Procedure

1. Open *Manage Databases* in the SAP HANA cockpit by drilling down from *Overall Tenant Statuses* in the system *Overview* for the system database.
2. Select *Configure Workload Allocation* from the overflow menu in the header.
3. Select *Memory Allocation*.
4. View the *Default allocation limit per service*, or choose to edit the default by clicking the pencil icon. Changing the default causes each of the databases and hosts to inherit the value of the allocation limit (except those databases or hosts that have an allocation limit that already differs from the default).
5. View the *Allocation Limit* for each database or host, or choose to edit the value by clicking the pencil.
6. If you choose to edit the limit value, the *Edit Allocation Limit* dialog displays. Enter a value, or clear the allocation limit input field to revert to the allocation limit inherited from the default value.

The total memory available for allocation is equal to the sum of the allocation limits for each database (regardless of whether you have edited the limit or allowed the default value to be inherited). A warning will display if this sum exceeds the global allocation limit (the amount of available memory per host). However, SAP HANA does allow the sum of the global allocation limits to exceed the global allocation limit.

Related Information

[SAP HANA Used Memory \[page 269\]](#)

[Memory Sizing \[page 271\]](#)

[Allocated Memory Pools and Allocation Limits \[page 272\]](#)

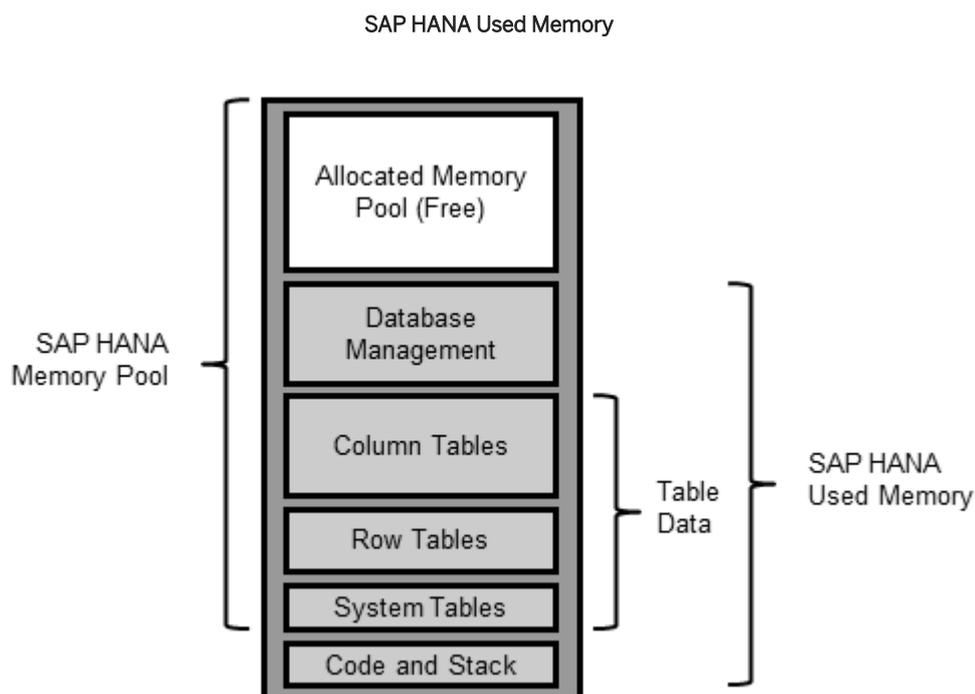
[SAP HANA Memory Usage and the Operating System \[page 273\]](#)

6.3.3.1.1 SAP HANA Used Memory

The total amount of memory used by SAP HANA is referred to as used memory. It includes program code and stack, all data and system tables, and the memory required for temporary computations.

SAP HANA consists of a number of processes running in the Linux operating environment. Under Linux, the operating system (OS) is responsible for reserving memory to all processes. When SAP HANA starts up, the OS reserves memory for the program code (sometimes called the text), the program stack, and static data. It then dynamically reserves additional data memory when requested by the SAP HANA memory manager. Dynamically allocated memory consists of heap memory and shared memory.

The following figure shows used memory, consisting of code, stack, and table data:



Since the code and program stack size are about 6 GB, almost all of used memory is used for table storage, computations, and database management.

Service Used Memory

An SAP HANA system consists of multiple services that all consume memory, in particular the `indexserver` service, the main database service. The index server holds all the data tables and temporary results, and therefore dominates SAP HANA used memory.

Peak Used Memory

Ultimately, it is more important to understand the behavior of used memory over time and under peak loads. For this purpose, SAP HANA has a special used memory indicator called peak used memory. As the value for used memory is a current measurement, peak used memory allows you to keep track of the maximum value for used memory over time.

You can also reset peak used memory. This can be useful if you want to establish the impact of a certain workload on memory usage. So for example, you can reset peak used memory, run the workload, and then examine the new peak used memory value.

Memory Usage of Tables

The dominant part of the used memory in the SAP HANA database is the space used by data tables. Separate measurements are available for column-store tables and row-store tables.

i Note

The SAP HANA database loads column-store tables into memory column by column only upon use. This is sometimes called "lazy loading". This means that columns that are never used will not be loaded and memory waste is avoided. When the SAP HANA database runs out of allocatable memory, it will try to free up some memory by unloading unimportant data (such as caches) and even table columns that have not been used recently. Therefore, if it is important to measure precisely the total, or worst-case, amount of memory used for a particular table, it is important to ensure that the table is first fully loaded into memory. You can do this by loading the table into memory.

Memory Usage of Expensive Statements

Every query and statement consumes memory, for the evaluation of the statement plan, caching, and, mainly the calculation of intermediate and final results. While many statement executions use only a moderate amount of memory, some queries, for instance using unfiltered cross joins, will tax even very large systems.

Expensive statements are individual SQL statements whose execution time exceeded a configured threshold. The expensive statements trace records information about these statements for further analysis. If in addition to activating the expensive statements trace, you enable per-statement memory tracking, the expensive statements trace will also show the peak memory size used to execute expensive statements.

It is further possible to protect an SAP HANA system against excessive memory usage due to uncontrolled queries by limiting the amount of memory used by single statement executions per host.

Related Information

[Monitoring and Analyzing with the Performance Monitor \[page 397\]](#)

[Monitor Tables by Size and Usage \[page 334\]](#)

[Load/Unload a Column Table into/from Memory \[page 520\]](#)

[Monitoring and Analyzing with the Statements Monitor \[page 406\]](#)

[Monitoring and Analyzing Expensive Statements \[page 407\]](#)

[Setting a Memory Limit for SQL Statements \[page 633\]](#)

6.3.3.1.2 Memory Sizing

Memory sizing is the process of estimating in advance the amount of memory that will be required to run a certain workload on an SAP HANA database. To understand memory sizing, several questions need to be answered.

- What is the size of the data tables that will be stored in the SAP HANA database?
You may be able to estimate this based on the size of your existing data, but unless you precisely know the compression ratio of the existing data and the anticipated growth factor, this estimate may not be accurate.
- What is the expected compression ratio that SAP HANA will apply to these tables?
The column store of the SAP HANA database automatically uses a combination of various advanced compression algorithms (dictionary, RLE, sparse, and so on) to compress each table column separately. The achieved compression ratio depends on many factors, such as the nature of the data, its organization and data types, the presence of repeated values, the number of indexes (SAP HANA requires fewer indexes), and so on.
- How much extra working memory will be required for temporary computations?
The amount of extra memory will depend on the size of the tables (larger tables will create larger intermediate result tables in operations such as joins), but even more on the expected workload in terms of the concurrency and complexity of analytical queries (each concurrent query needs its own workspace).

The following SAP Notes provide additional tools and information to help you size the required amount of memory:

- SAP Note 1514966 - SAP HANA 1.0: Sizing SAP In-Memory Database
- SAP Note 1637145 - SAP BW on HANA: Sizing SAP In-Memory Database
- SAP Note 2296290 - New Sizing Report for BW on HANA

However, the most accurate method is to import several representative tables into an SAP HANA system, measure the memory requirements, and extrapolate from the results.

Related Information

[SAP Note 1514966](#)

[SAP Note 1637145](#)

[SAP Note 2296290](#)

6.3.3.1.3 Allocated Memory Pools and Allocation Limits

SAP HANA, across its different processes, reserves a pool of memory before actual use. This pool of allocated memory is preallocated from the operating system over time, up to a predefined global allocation limit, and is then efficiently used by SAP HANA as needed.

SAP HANA preallocates and manages its own memory pool, used for storing in-memory table data, thread stacks, temporary results, and other system data structures. When more memory is required for table growth or temporary computations, the SAP HANA memory manager obtains it from the pool. When the pool cannot satisfy the request, the memory manager increases the pool size by requesting more memory from the operating system, up to a predefined allocation limit.

By default, the allocation limit is calculated as follows: 90% of the first 64 GB of available physical memory on the host plus 97% of each further GB.

There is normally no reason to change the value of this variable, unless you purchased a license for less than the total amount of physical memory. In this case, you need to change the global allocation limit to remain in compliance with the license.

❖ Example

- You have a server with 512GB, but purchased an SAP HANA license for only 384 GB. You therefore set the `global_allocation_limit` to 393216 (384 * 1024 MB).
- You have a distributed HANA system on four hosts with 512 GB each, but purchased an SAP HANA license for only 768 GB. Set the `global_allocation_limit` to 196608 (192 * 1024 MB on each host).

Another case in which you may want to limit the size of the memory pool is on development systems with more than one SAP HANA system installed on a single host. This will avoid resource contentions or conflicts.

Service Allocation Limit

In addition to the global allocation limit, each service running on the host has an allocation limit, the service allocation limit. Given that collectively, all services cannot consume more memory than the global allocation limit, each service has what is called an effective allocation limit. The effective allocation limit of a service specifies how much physical memory a service can in reality consume given the current memory consumption of other services.

❖ Example

A single-host system has 100 GB physical memory. Both the global allocation limit and the individual service allocation limits are 92.5% (default values). This means the following:

- Collectively, all services of the SAP HANA database can use a maximum of 92.5 GB.
- Individually, each service can use a maximum of 92.5 GB.

Therefore, if 2 services are running and the current memory pool of service 1 is 50 GB, then the effective allocation limit of service 2 is 42.5 GB. This is because service 1 is already using 50 GB and together they cannot exceed the global allocation limit of 92.5 GB.

What happens when the allocation limit is reached?

Memory is a finite resource. Once the allocation limit has been reached and the pool is exhausted, the memory manager can no longer allocate memory for internal operations without first giving up something else. Buffers and caches are released, and column store tables are unloaded, column by column, based on a least-recently-used order, up to a preset lower limit. When tables are partitioned over several hosts, this is managed on a host-by-host basis; that is, column partitions are unloaded only on hosts with an acute memory shortage.

Table (column or partition) unloading is generally not a good situation since it leads to performance degradation later when the data will have to be reloaded for queries that need them. You can identify pool exhaustion by examining the `M_CS_UNLOADS` system view.

However, it is still possible that the memory manager needs more memory than is available. For example, when too many concurrent transactions use up all memory, or when a particularly complex query performs a cross join on very large tables and creates a huge intermediate result that exceeds the available memory. Such situations can potentially lead to an out-of-memory failure.

Related Information

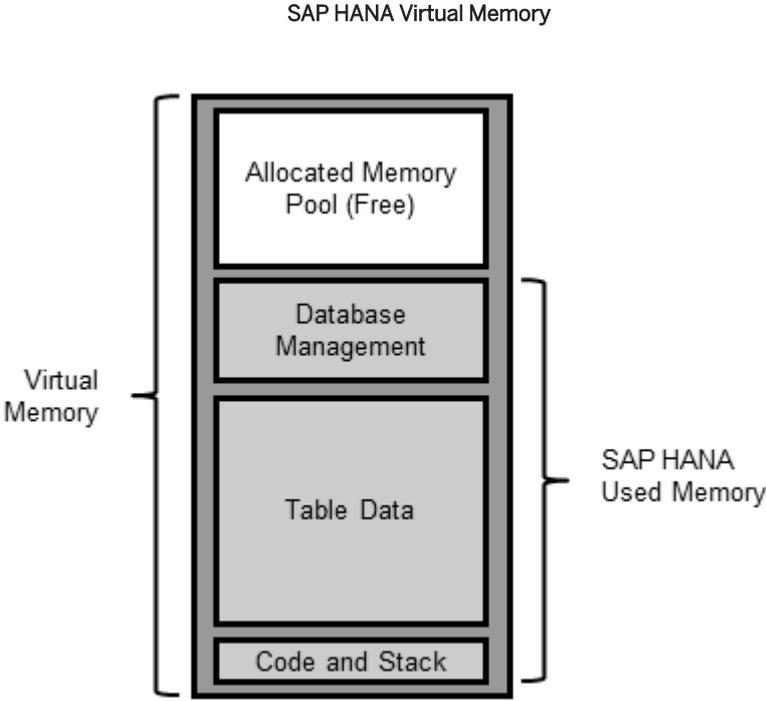
[Change the Global Memory Allocation Limit \[page 478\]](#)

6.3.3.1.4 SAP HANA Memory Usage and the Operating System

Due to the way in which SAP HANA manages memory, the relationship between Linux memory indicators and SAP HANA's own memory indicators may not correlate as expected.

From the perspective of the Linux operating system, SAP HANA is a collection of separate processes. Linux programs reserve memory for their use from the Linux operating system. The entire reserved memory footprint of a program is referred to as its virtual memory. Each Linux process has its own virtual memory, which grows when the process requests more memory from the operating system, and shrinks when the process relinquishes unused memory. You can think of virtual memory size as the memory amount that the

process has requested (or allocated) from the operating system, including reservations for its code, stack, data, and memory pools under program control. SAP HANA's virtual memory is logically shown in the following figure:



i Note
 SAP HANA really consists of several separate processes, so the figure above shows all SAP HANA processes combined.

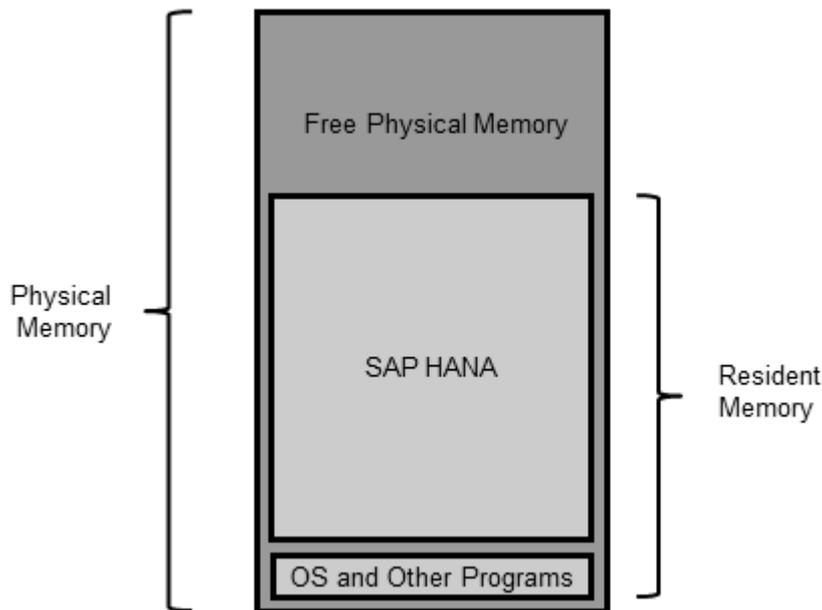
Virtual, Physical, and Resident Memory

When part of the virtually allocated memory actually needs to be used, it is loaded or mapped to the real, physical memory of the host and becomes resident. Physical memory is the DRAM memory installed on the host. On most SAP HANA hosts, it ranges from 256 gigabytes (GB) to 1 terabyte (TB). It is used to run the Linux operating system, SAP HANA, and all other programs.

Resident memory is the physical memory actually in operational use by a process. Over time, the operating system may swap out some of a process's resident memory according to a least-recently-used algorithm to make room for other code or data. Thus, a process's resident memory size may fluctuate independently of its virtual memory size. In a properly-sized SAP HANA appliance, there is enough physical memory, so that swapping is disabled and should not be observed.

This can be illustrated as follows:

SAP HANA Resident Memory



On a typical SAP HANA appliance, the resident memory part of the operating system and all other running programs usually does not exceed 2 GB. The rest of the memory is therefore dedicated for the use of SAP HANA.

When memory is required for table growth or for temporary computations, the SAP HANA code obtains it from the existing memory pool. When the pool cannot satisfy the request, the SAP HANA memory manager will request and reserve more memory from the operating system. At this point, the virtual memory size of SAP HANA processes grows.

Once a temporary computation completes or a table is dropped, the freed memory is returned to the memory manager, which recycles it to its pool without informing the operating system. Therefore, from SAP HANA's perspective, the amount of used memory shrinks, but the processes' virtual and resident memory sizes are not affected. This creates a situation where the used memory value may shrink to below the size of SAP HANA's resident memory. This is normal.

i Note

The memory manager may also choose to return memory back to the operating system, for example when the pool is close to the allocation limit and contains large unused parts.

Related Information

[SAP HANA Used Memory \[page 269\]](#)

[Memory Sizing \[page 271\]](#)

[Allocated Memory Pools and Allocation Limits \[page 272\]](#)

6.3.3.2 Define CPU Cores Allocation Limits

As part of the provisioning process, you can ensure that CPU cores are shared appropriately between tenant databases. By setting the allocation limits, you ensure appropriate sharing of CPU cores, effectively specifying the maximum number of CPU cores that can be used by a particular tenant on a particular host.

Context

In the SAP HANA cockpit you can use the *CPU Allocation* tab of the *Configure Workload Allocation* app to view and modify the CPU cores allocation limit. This limit corresponds to the settings of the `max_concurrency` parameter in the `global.ini` file.

Each host in an SAP HANA system has a physical set of CPU cores. You can adjust the maximum number of CPU cores available for allocation at each of the following levels:

- System
- Tenant database (these settings override the system settings)
- Host (these settings override the database and system settings)

For each database (or each database per host in a multi-host system), you can also refer to a mini-chart representing the CPU usage of the indexserver in one day, where dark green shows the peak CPU usage, and light green shows the average CPU usage.

Procedure

1. Open *Manage Databases* in the SAP HANA cockpit by drilling down from *Overall Tenant Statuses* in the system *Overview* for the system database.
2. Select *Configure Workload Allocation* from the overflow menu in the header.
3. Select *CPU Allocation*.
4. View the *Default CPU cores allocation limit*, or choose to edit the default by clicking the pencil icon. Changing the default causes each of the databases and hosts to inherit the value of the allocation limit (except those databases or hosts that have an allocation limit that already differs from the default).
5. View the *CPU Core Limit* for each of the databases, or choose to edit the value by clicking the pencil.
6. If you choose to edit the limit value, the *Edit Max CPU Cores* dialog displays. Enter a value, or clear the allocation limit input field to revert to the allocation limit inherited from the default value.

The total number of CPU cores available for allocation is equal to the sum of the allocation limits for each database (regardless of whether you have edited the limit or allowed the default value to be inherited). A warning will display if this sum exceeds the number of physical CPU cores.

Related Information

[Controlling CPU Consumption \[page 277\]](#)

6.3.3.2.1 Controlling CPU Consumption

If the physical hardware on a host is shared between several processes you can use CPU affinity settings to assign a set of logical cores to a specific SAP HANA process. These settings are coarse-grained and apply on the OS and process-level.

Prerequisites

Using this workload management option, we firstly analyze how the system CPUs are configured and then, based on the information returned, apply affinity settings in `daemon.ini` to bind specific processes to logical CPU cores. Processes must be restarted before the changes become effective. This approach applies primarily to the use cases of SAP HANA tenant databases and multiple SAP HANA instances on one server; you can use this, for example, to partition the CPU resources of the system by tenant database.

→ Tip

As an alternative to applying CPU affinity settings you can achieve similar performance gains by changing the parameter `[execution] max_concurrency` in the `global.ini` configuration file. This may be more convenient and does not require the system to be offline.

To make the changes described here you require access to the operating system of the SAP HANA instance to run the Linux `lscpu` command and you require the privilege INIFILE ADMIN.

Information about the SAP HANA system topology is also available from SAP HANA monitoring views as described in the following subsection *SAP HANA Monitoring Views for CPU Topology Details*.

Context

For Xen and VMware, the users in the VM guest system see what is configured in the VM host. So the quality of the reported information depends on the configuration of the VM guest. Therefore SAP cannot give any performance guarantees in this case.

Procedure

1. Firstly, to confirm the physical and logical details of your CPU architecture, analyze the system using the `lscpu` command. This command returns a listing of details of the system architecture. The table which follows gives a commentary on the most useful values based on an example system with 2 physical chips (sockets) each containing 8 physical cores. These are hyperthreaded to give a total of 32 logical cores.

#	Feature	Example Value
1	Architecture	x86_64
2	CPU op-mode(s)	32-bit, 64-bit
3	Byte Order	LittleEndian
4	CPUs	32
5	On-line CPU(s) list	0-31
6	Thread(s) per core	2
7	Core(s) per socket	8
8	Socket(s)	2
9	NUMA node(s)	2
21	NUMA node0 CPU(s)	0-7,16-23
22	NUMA node1 CPU(s)	8-15,24-31

- 4-5: This example server has 32 logical cores numbered 0 - 31
- 6-8: Logical cores ("threads") are assigned to physical cores. Where multiple threads are assigned to a single physical core this is referred to as 'hyperthreading'. In this example, there are 2 sockets, each socket contains 8 physical cores (total 16). Two logical cores are assigned to each physical core, thus, each core exposes two execution contexts for the independent and concurrent execution of two threads.
- 9: In this example there are 2 NUMA (Non-uniform memory access) nodes, one for each socket. Other systems may have multiple NUMA nodes per socket.
- 21-22: The 32 logical cores are numbered and specifically assigned to one of the two NUMA nodes.

i Note

Even on a system with 32 logical cores and two sockets the assignment of logical cores to physical CPUs and sockets can be different. It is important to collect the assignment in advance before making changes. A more detailed analysis is possible using the system commands described in the next step. These provide detailed information for each core including how CPU cores are grouped as siblings.

2. In addition to the `lscpu` command you can use the set of system commands in the `/sys/devices/system/cpu/` directory tree. For each logical core there is a numbered subdirectory beneath this node (`/cpu12/` in the following examples). The examples show how to retrieve this information and the table gives details of some of the most useful commands available:

❖ Example

```
cat /sys/devices/system/cpu/present
cat /sys/devices/system/cpu/cpu12/topology/thread_siblings_list
```

Command	Example Output	Commentary
present	0-15	The number of logical cores available for scheduling.
cpu12/topology/core_siblings_list	4-7, 12-15	The cores on the same socket.
cpu12/topology/thread_siblings_list	4, 12	The logical cores assigned to the same physical core (hyperthreading).
cpu12/topology/physical_package_id	1	The socket of the current core - in this case cpu12.

- Based on the results returned you can use the `affinity` setting to restrict CPU usage of SAP HANA processes to certain CPUs or ranges of CPUs. You can do this for the following servers: nameserver, indexserver, compileserver, preprocessor, and xsengine (each server has a section in the `daemon.ini` file). The examples and commentary below show the syntax for the ALTER SYSTEM CONFIGURATION commands required. The changed affinity settings only take effect after a restart of the affected SAP HANA processes.

❁ Example

To restrict the nameserver to two logical cores of the first CPU of socket 0 (see line 21 in the example above), use the following affinity setting:

```
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('nameserver', 'affinity') = '0,16'
```

❁ Example

To restrict the preprocessor and the compileserver to all remaining cores (that is, all except 0 and 16) on socket 0 (see line 21 in the example above), use the following affinity settings:

```
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('preprocessor', 'affinity') = '1-7,17-23'
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('compileserver', 'affinity') = '1-7,17-23'
```

❁ Example

To restrict the indexserver to all cores on socket 1 (see line 22 in the example above), use the following affinity settings:

```
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('indexserver', 'affinity') = '8-15,24-31'
```

❁ Example

To set the affinity for two tenant databases called DB1 and DB2 respectively in a tenant database setup, use the following affinity settings:

```
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('indexserver.DB1', 'affinity') = '1-7,17-23';
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('indexserver.DB2', 'affinity') = '9-15,25-31';
```

Other Linux commands which are relevant here are `sched_setaffinity` and `numactl`:
`sched_setaffinity` limits the set of CPU cores available (by applying a CPU affinity mask) for execution of a specific process (this could be used, for example, to isolate tenants) and `numactl` controls NUMA policy for processes or shared memory.

Related Information

[Configuring Memory and CPU Usage for Tenant Databases \[page 266\]](#)

[SAP HANA Monitoring Views for CPU Topology Details \[page 630\]](#)

6.3.4 Using SAP Web Dispatcher for Load Balancing with Tenant Databases

If an SAP HANA system has multiple instances of SAP HANA extended services, classic model (SAP HANA XS classic) and is distributed across multiple hosts, you can implement an external SAP Web Dispatcher to distribute the load of inbound HTTP requests and to ensure high availability.

The following section describes how to configure an external SAP Web Dispatcher for SAP HANA systems with tenant databases.

Before You Start

Note the following points:

- The external SAP Web Dispatcher is a separate installation and does not form part of the SAP HANA system. It must have a minimum version of **745 Patch Level 21**.
- An SAP Web Dispatcher process also runs on all SAP HANA hosts on which an instance of SAP HANA XS is active. This internal SAP Web Dispatcher is a fixed part of the SAP HANA system. In a system with tenant databases, this internal SAP Web Dispatcher must also be configured to enable HTTP access to individual databases. For more information, see *Configure HTTP(S) Access to Tenant Databases*.
- All information and configuration steps described in SAP Note [1855097](#) are still valid. In particular, the parameter `wdisp/filter_xs_internal_uri` has to be set to `false` in the `webdispatcher.ini` configuration file of your SAP HANA system.
- The configuration described in the following sections describes access to tenant databases. However, it is also valid for the system database. For the Web Dispatcher, there is no difference between tenant databases and the system database.
- The SAP Web Dispatcher handles only HTTP(S) access to SAP HANA.
- For more information about configuring secure HTTPS access, see *Configure HTTP(S) Access to Tenant Databases* (internal Web Dispatcher configuration) and *Configuring SAP Web Dispatcher to Support SSL* in the SAP HANA Web Dispatcher documentation.

Related Information

[Configure HTTP\(S\) Access to Tenant Databases via SAP HANA XS Classic \[page 1578\]](#)

[Configuring SAP Web Dispatcher to Support SSL](#)

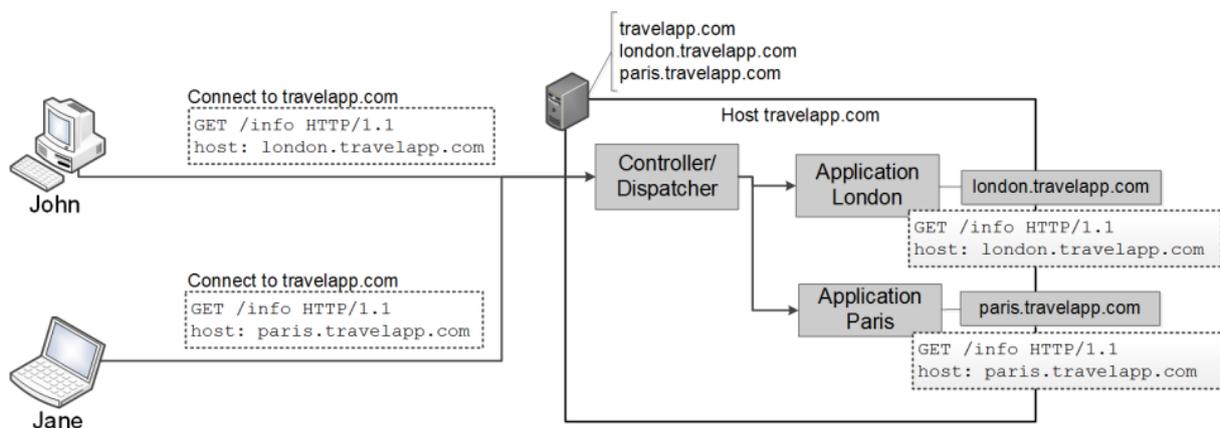
[Virtual-Host-Based Routing \[page 281\]](#)

[Configuring an External SAP Web Dispatcher for Tenant Databases \[page 283\]](#)

6.3.4.1 Virtual-Host-Based Routing

An example explains the basics of virtual-host-based routing.

The Website `travelapp.com` provides Web-based services for information about popular travel destinations. Services are implemented as separate applications, which run on separate Web servers on one host (`travelapp.com`). Virtual host names are used to distinguish between the available services: `london.travelapp.com` and `paris.travelapp.com`. Both virtual host names are aliases for `travelapp.com`. This can be illustrated as follows:



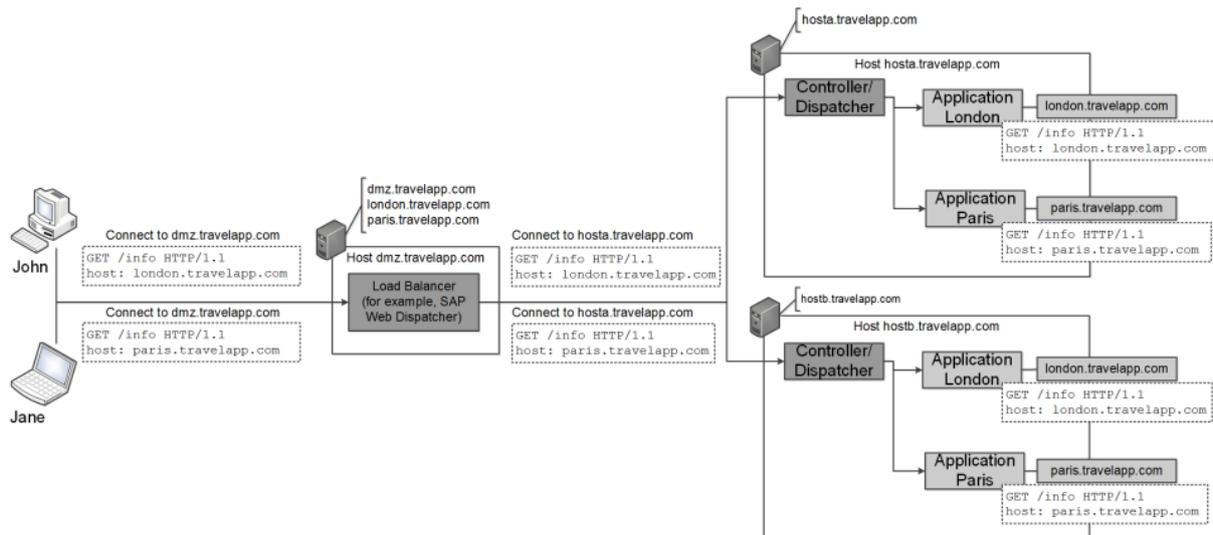
Virtual-Host-Based Routing

John wants to read information about London. Therefore, he enters `london.travelapp.com` into his browser. As `london.travelapp.com` is an alias for `travelapp.com`, the browser sends the HTTP request to `travelapp.com`, but it uses `london.travelapp.com` as the host header of this request. The request arrives at a controller or dispatcher process on `travelapp.com`. This dispatcher process decides whether to forward the request to the Web server responsible for displaying information about London or the Web server responsible for displaying information about Paris. This decision is made based on the host header, that is the host name that the user originally entered into the browser. `london.travelapp.com` is assigned to the application for London and `paris.travelapp.com` is assigned to the application for Paris.

Jane requires information about Paris and enters `paris.travelapp.com` into her browser. This request also arrives at the dispatcher process and is dispatched to the Paris application based on the host header of the request.

Load Balancing

travelapp.com has proved to be a successful service with many users. As a result, one host is no longer sufficient, and the application has been installed on a second host. In addition, a load balancer is needed to distribute requests between the two hosts. The aliases london.travelapp.com and paris.travelapp.com have to be changed to point to the host of the load balancer to guarantee that all requests are handled by the load balancer (dmz.travelapp.com). This can be illustrated as follows:

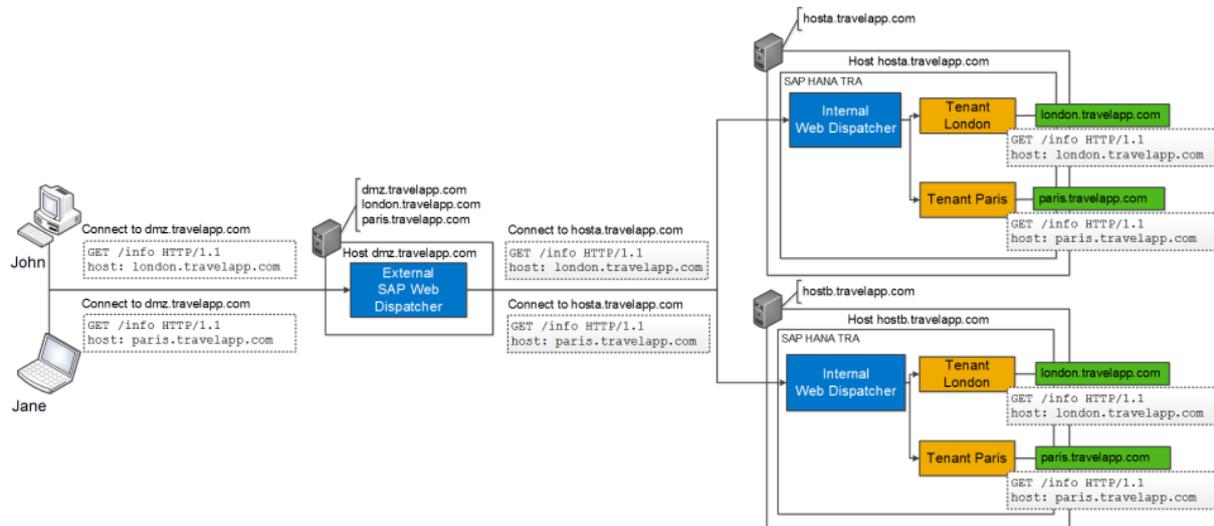


Virtual-Host-Based Routing with Load Balancing

John again wants to read information about London. Therefore, he enters london.travelapp.com into his browser. As london.travelapp.com is an alias for dmz.travelapp.com, the browser sends the HTTP request to dmz.travelapp.com, but it uses london.travelapp.com as the host header of this request. This request arrives at the load balancer, which simply forwards the request to hosta.travelapp.com or hostb.travelapp.com based on the current load. It must not change the host header of the request because this request is later necessary in the dispatcher. After that, the dispatcher handles the request as if no load balancer is involved, regardless of the fact that the host name in the host header actually points to another host.

SAP HANA Tenant Databases

Translated to the context of SAP HANA tenant databases, the load balancer is an external SAP Web Dispatcher, and the dispatcher is the system-internal SAP Web Dispatcher, as illustrated in the following figure:



Virtual-Host-Based Routing for SAP HANA Tenant Databases

6.3.4.2 Configuring an External SAP Web Dispatcher for Tenant Databases

Virtual host names for differentiated HTTP access to tenant databases are configured in the system-internal SAP Web Dispatcher. If you're using an external SAP Web Dispatcher for load balancing, you must also configure the external Web Dispatcher. Otherwise, information about the selected virtual hosts can't be transported to the SAP HANA system.

→ Remember

All of the configuration settings mentioned here are done in the **external** Web Dispatcher and not in the internal Web Dispatcher that is part of the SAP HANA system. The external Web Dispatcher is a separate installation and does not form part of the SAP HANA system. Before you can configure the external Web Dispatcher, the internal Web Dispatcher must already have been configured to enable HTTP access to individual databases. For more information, see *Configure HTTP(S) Access to Tenant Databases*.

Single Versus Multiple Tenant Access via External Web Dispatcher

Every tenant database that needs to be accessed through the external Web Dispatcher requires a `wdisp/system_<xx>` parameter entry in the external Web Dispatcher profile (`sapwebdisp.pfl`). The `XSSRV` subparameter specifies the XS server to connect to, and the `XSVHOST` subparameter specifies the virtual host name of the tenant database.

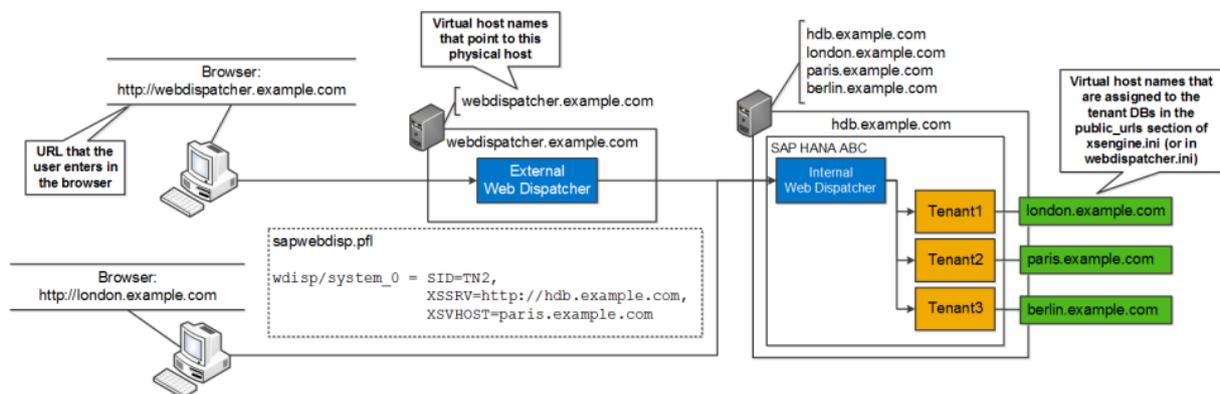
❁ Example

```
wdisp/system_<xx> = SID=<3-digit ID>, XSSRV=http://<physical host name of SAP HANA server>:<port>, XSVHOST=<virtual host name of the tenant>
```

i Note

Virtual host names are configured in the `public_urls` section of the `xsengine.ini` configuration file (or in `webdispatcher.ini`). This is part of the **internal** Web Dispatcher configuration. For more information, see *Configure HTTP(S) Access to Tenant Databases*.

If only **one tenant database** needs to be accessed through the external Web Dispatcher, a single `wdisp/system_<XX>` entry for the tenant database with the above configuration is sufficient, as depicted in the following figure:



i Note

The figure shows a simplified depiction of the Web Dispatcher profile (`sapwebdisp.pfl`). In the real configuration, the `XSSRV` subparameter requires port numbers, and line breaks are not allowed.

Access to a Single Tenant Database

i Note

An external Web Dispatcher is not mandatory to access a single tenant. But there are scenarios in which an external Web Dispatcher is required, for example sophisticated applications (for example, some SAP Fiori scenarios) or for security purposes. For more information, see the relevant application documentation and the SAP Web Dispatcher documentation.

Virtual host names are used to configure tenant differentiation. Two scenarios are possible:

- **Option 1:** Tenant databases are accessed via HTTP through the external Web Dispatcher only; there is no direct HTTP access to the tenant databases (recommended)
- **Option 2:** Tenants databases are accessed via HTTP both through the external Web Dispatcher and directly, bypassing the external Web Dispatcher
This configuration requires additional virtual host names and is more complex than option 1. However, this option is useful if the external Web Dispatcher is being added to an existing landscape.

Related Information

[SAP Web Dispatcher Documentation on SAP Help Portal](#)

[Configure HTTP\(S\) Access to Tenant Databases via SAP HANA XS Classic \[page 1578\]](#)

[Option 1: Configuring Access to Multiple \(or All\) Tenant Databases Through External Web Dispatcher Only \[page 285\]](#)

[Option 2: Configuring Access to Multiple \(or All\) Tenant Databases Through External Web Dispatcher and Directly \[page 288\]](#)

6.3.4.2.1 Option 1: Configuring Access to Multiple (or All) Tenant Databases Through External Web Dispatcher Only

Use this configuration if you want tenant databases to be accessed through the external Web Dispatcher only. With this configuration, there is no direct HTTP access to the tenant databases.

The main part of this configuration involves setting the virtual host names of tenant databases configured in the external Web Dispatcher profile to point to the host of the external Web Dispatcher, instead of the host of the SAP HANA system. As a result, all requests to the virtual host name of a tenant database first go to the external Web Dispatcher and are then forwarded to the internal Web Dispatcher in the SAP HANA system.

→ Remember

Before you can configure the external Web Dispatcher, the internal Web Dispatcher must already have been configured to enable HTTP access to individual databases. For more information, see *Configure HTTP(S) Access to Tenant Databases*.

Single-Host Systems

The `wdisp/system_<xx>` entry for each tenant database is configured in the external Web Dispatcher profile as follows:

- `XSSRV` specifies the actual physical SAP HANA host name and port to which requests are sent.
- `XSVHOST` specifies the virtual host name of the tenant database to which requests are sent.

i Note

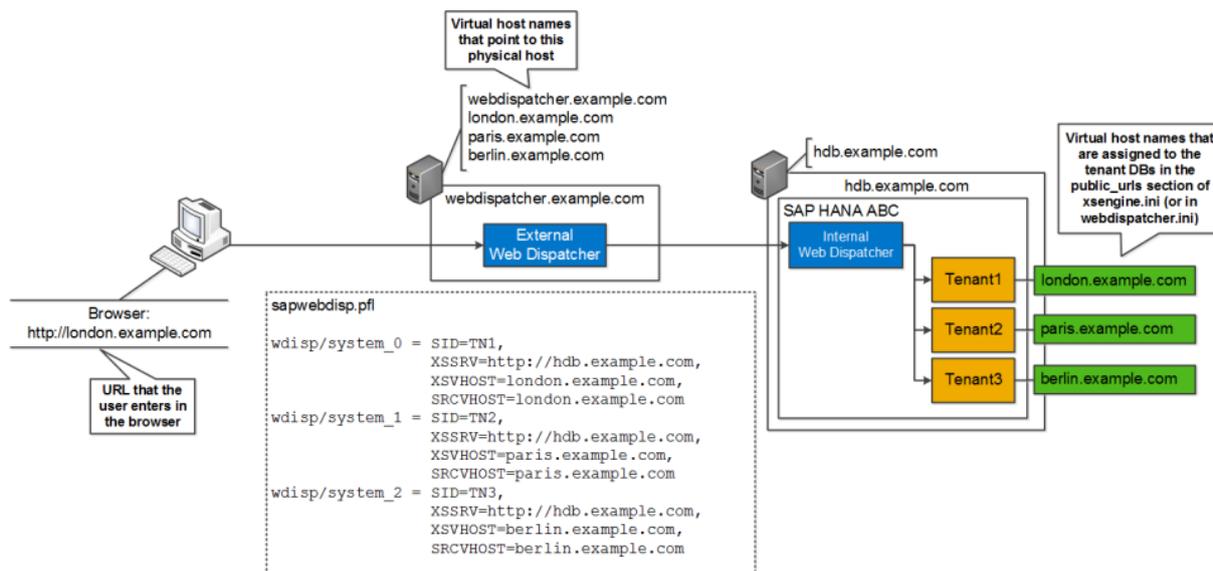
If a tenant database has multiple virtual host names assigned, only one needs to be entered in `XSVHOST`.

- `SRCVHOST` specifies the virtual host name that is used to map incoming HTTP requests to the `wdisp/system` entry that represents a particular tenant.

i Note

With this configuration option, `XSVHOST` and `SRCVHOST` are always identical.

The following figure depicts this configuration in a single-host system:



Access to Multiple Tenant Databases (Single Host)

Note

The figure shows a simplified depiction of the Web Dispatcher profile (`sapwebdisp.pfl`). In the real configuration, the `XSSRV` subparameter requires port numbers, and line breaks are not allowed.

In the example depicted above, what happens when the user enters `london.example.com` into her browser?

1. The browser opens a TCP/IP connection to `webdispatcher.example.com` because `london.example.com` is only an alias name for `webdispatcher.example.com`.
2. The browser sends an HTTP request over this connection. The host header of this HTTP request is `london.example.com`, which is the URL that the user entered.
3. The external Web Dispatcher receives the HTTP request, checks the host header and uses this to map the request to a `wdisp/system` entry. As `london.example.com` is the `SRCVHOST` value for `wdisp/system_0`, the request is associated with `wdisp/system_0`.
4. The external Web Dispatcher opens a TCP/IP connection to the `XSSRV` value of `wdisp/system_0` (`hdb.example.com`).
5. The external Web Dispatcher sets the destination of the request to the tenant database specified in the `XSVHOST` subparameter of `wdisp/system_0` (`london.example.com`) by injecting a proprietary HTTP header into the request.
6. The internal SAP HANA Web Dispatcher receives the request. Because of the injected HTTP header field, it identifies that the request is destined for tenant database 1 and forwards it to the XS server of tenant database 1.

Multiple-Host Systems

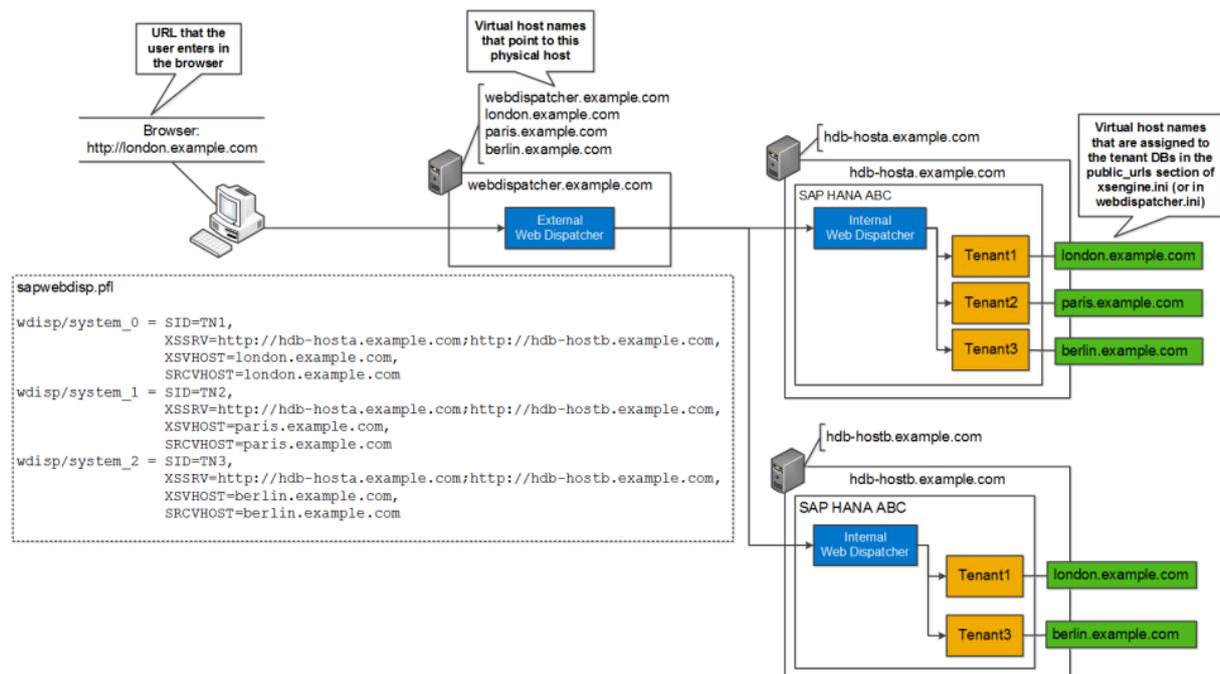
In a multiple-host system, the external Web Dispatcher must be configured to connect to all hosts. This means that all hosts with a running XS server (or that may have an XS server in the future) have to be entered as the

value for `XSSRV` as a semi-colon (;) separated list. Even if a tenant database is not running on a host, you should add the host to the list anyway. This will enable the smooth moving of tenant databases without the need to change the external Web Dispatcher configuration.

→ Remember

In the internal Web Dispatcher configuration, the virtual host name in the `XSVHOST` subparameter must be assigned to the tenant database on **all** hosts.

The following figure depicts the configuration in a multiple-host system:



i Note

The figure shows a simplified depiction of the Web Dispatcher profile (`sapwebdisp.pfl`). In the real configuration, the `XSSRV` subparameter requires port numbers, and line breaks are not allowed.

Access to Multiple Tenant Databases (Multiple Hosts)

Now what happens when the user enters `london.example.com` into her browser?

The process is identical to the single-host scenario with one exception: The external Web Dispatcher periodically checks which host(s) a tenant database is actually running on. If a tenant database is running on multiple hosts, the external Web Dispatcher performs load balancing between these hosts.

Related Information

[Configure HTTP\(S\) Access to Tenant Databases via SAP HANA XS Classic \[page 1578\]](#)

6.3.4.2.2 Option 2: Configuring Access to Multiple (or All) Tenant Databases Through External Web Dispatcher and Directly

Use this configuration if you want tenant databases to be accessed both through the external Web Dispatcher and directly, bypassing the external Web Dispatcher.

With this configuration, additional virtual host names are required for each tenant database. These virtual host names point to the physical host name of the external Web Dispatcher. The virtual host names that are assigned to the tenant databases still point to the host of the SAP HANA system.

→ Remember

Before you can configure the external Web Dispatcher, the internal Web Dispatcher must already have been configured to enable HTTP access to individual databases. For more information, see *Configure HTTP(S) Access to Tenant Databases*.

Single-Host Systems

The `wdisp/system_<xx>` entry for each tenant database is then configured in the external Web Dispatcher profile as follows:

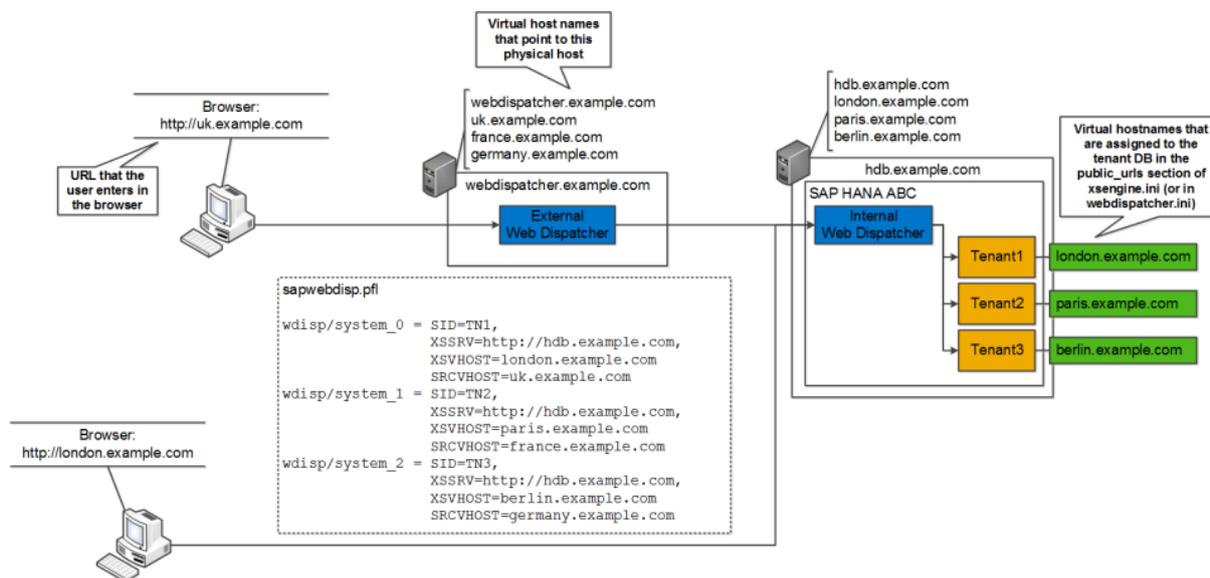
- `XSSRV` specifies the actual physical host name and port to which requests are sent.
- `XSVHOST` specifies the virtual host name of the tenant database to which requests are sent.

i Note

If a tenant database has multiple virtual host names assigned, only one needs to be entered in `XSVHOST`.

- `SRCVHOST` specifies the virtual host name that is used to map incoming HTTP requests to the `wdisp/system_<xx>` that represents a particular tenant.

The following figure depicts this configuration in a single-host system:



Access to Multiple Tenant Databases in Single-Host System

i Note

The figure shows a simplified depiction of the Web Dispatcher profile (`sapwebdisp.pfl`). In the real configuration, the `XSSRV` subparameter requires port numbers, and line breaks are not allowed.

In the example depicted above, what happens when the user enters `uk.example.com` into his browser?

1. The browser opens a TCP/IP connection to `webdispatcher.example.com` because `uk.example.com` is only an alias name for `webdispatcher.example.com`.
2. The browser sends an HTTP request over this connection. The host header of this HTTP request is `uk.example.com`, which is the URL that the user entered.
3. The external Web Dispatcher receives the HTTP request, checks the host header and uses it to map the request to a `wdisp/system` entry. As `uk.example.com` is the `SRCVHOST` value for `wdisp/system_0`, the request is associated with `wdisp/system_0`.
4. The external Web Dispatcher opens a TCP/IP connection to the `XSSRV` value of `wdisp/system_0` (`hdb.example.com`).
5. The external Web Dispatcher sets the destination of the request to the tenant database specified in the `XSVHOST` parameter of `wdisp/system_0` (`london.example.com`) by injecting a proprietary HTTP header into the request.
6. The internal SAP HANA Web Dispatcher receives the request. Because of the injected HTTP header field, it identifies that the request is destined for tenant database 1 and forwards it to the XS server of tenant database 1.

Multiple-Host Systems

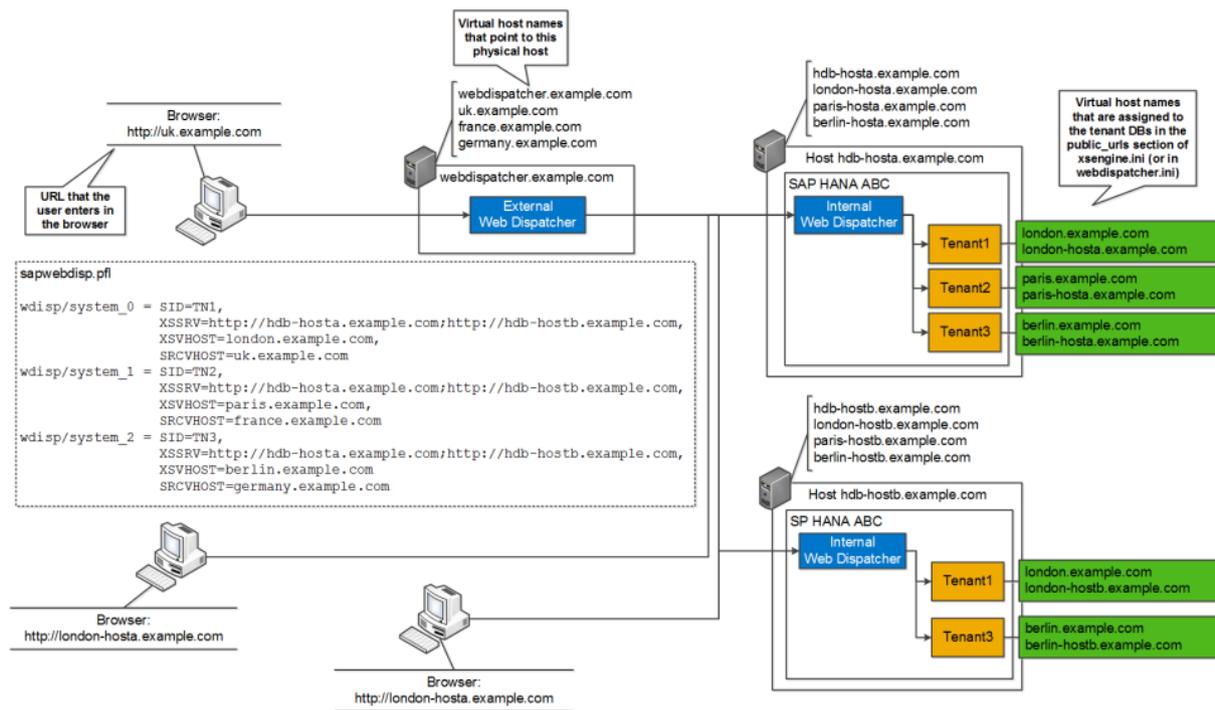
In a multiple-host system, the external Web Dispatcher must be configured to connect to all hosts. This means that all hosts with a running XS server (or that may have an XS server in the future) have to be entered as the

value for `XSSRV` as a semi-colon (;) separated list. Even if a tenant database is not running on a host, you should add the host to the list anyway. This will enable the smooth moving of tenant databases without the need to change the external Web Dispatcher configuration.

→ Remember

In the internal Web Dispatcher configuration, the virtual host name in the `XSVHOST` subparameter must be assigned to the tenant on **all** hosts.

The following figure depicts the configuration in a multiple-host system:



i Note

The figure shows a simplified depiction of the Web Dispatcher profile (`sapwebdisp.pfl`). In the real configuration, the `XSSRV` subparameter requires port numbers, and line breaks are not allowed.

Access to Multiple Tenant Databases in Multiple-Host System

What happens after the user enters `uk.example.com` into his browser?

The process is identical to the single-host scenario with one exception: The external Web Dispatcher periodically checks on which host(s) a tenant database is actually running. If a tenant database is running on multiple hosts, the external Web Dispatcher performs load balancing between these hosts.

Related Information

[Configure HTTP\(S\) Access to Tenant Databases via SAP HANA XS Classic \[page 1578\]](#)

6.4 Configuring SAP HANA System Properties (INI Files)

An SAP HANA system has several configuration (*.ini) files that contain properties for configuring the system as a whole and individual tenant databases, as well as hosts and services.

About Configuration Files

Every SAP HANA service has its own configuration file (for example, `indexserver.ini`, `compileserver.ini`, and so on). For service-independent properties, there is also a `global.ini`.

Configuration files are separated into sections; sections bundle properties of the same category.

SAP HANA uses a layered configuration framework. This means that properties can be configured at different levels or layers depending on the configuration file. The following layers are available:

Layer	Description
Default	The default value for the property
System	The system-specific value for the property (configurable in the system database) If a system-specific value is not configured for a property, the default value applies.
Database	The database-specific value for the property (configurable in the system or tenant database) For some properties, it is possible to set database-specific values. If a database-specific value is not configured, the system-specific or host-specific value applies. For more information about layered configuration, see the section on database-specific configuration parameters.
Host	The host-specific value for the property (configurable in the system database) For some properties, it is possible to set host-specific values for multiple-host systems. If a host-specific value is not configured for a property that can be set at host level, the system-specific value applies.

Configuration files are stored on the SAP HANA server at the following locations according to layer:

- Default: `/usr/sap/<SID>/HDB<instance>/exe/config` (read only)
- System: `<sapmnt>/<SID>/SYS/global/hdb/custom/config`
- Database: `<sapmnt>/<SID>/SYS/global/hdb/custom/config/DB_<dbname>`
- Host: `/usr/sap/<SID>/HDB<instance>/<hostname>`

i Note

By default, `<sapmnt>` is `/hana/shared`.

The system view `M_INIFILES` contains information about the layers on which the properties of each configuration file can be configured. The system view `M_INIFILE_CONTENTS` contains information about the actual values configured for the properties of each file and on which layers.

Changing Parameter Values

You can change configuration parameters using the SAP HANA cockpit, the SAP HANA studio, or the ALTER SYSTEM ALTER CONFIGURATION statement.

Note

You cannot alter parameters of the secondary site in a system replication scenario using the SAP HANA cockpit or studio. Instead you must alter the parameters directly in the *.ini files on the secondary site. Afterwards reconfigure the database using `hdbnsutil -reconfig`.

In general, we do not recommend that you change the default values of parameters unless stated in the documentation or instructed to do so by SAP Support. For more information about frequently used parameters, see SAP Note 2036111.

While most parameters can be changed when the database is running, changes to some parameters require a database restart to take effect. To find out whether or not a restart is required for frequently used parameters, refer to SAP Note 2036111. Alternatively, you can see this information for each parameter by viewing the *.ini file in the file system.

Parameter Tracking

For traceability purposes changes to configuration values can be logged; optionally, a reason for the change can be entered in a comment value.

The tracking feature is enabled by default, but can be disabled by setting configuration parameter `write_log` in the `indexserver.ini` file to **false**. When it is enabled the SQL ALTER CONFIGURATION statement automatically updates the view `SYS.M_INIFILE_CONTENT_HISTORY`, this includes details of the time the change was made, the user name, the current and previous values and any comment text that was entered.

The following SQL example shows how the COMMENT keyword is used to include a reason for the change (refer to the *SAP HANA SQL and System Views Reference* for more details):

```
ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini', 'DATABASE', 'C11')
SET ('memorymanager', 'allocationlimit') = '500000'
WITH RECONFIGURE
COMMENT 'Reverting to previous setting';
```

An SQL command CLEAR is also available to truncate the ini file history:

```
ALTER SYSTEM CLEAR INIFILE CONTENT HISTORY [ UNTIL <timestamp> ]
```

This automatically adds a timestamp and comment to the history to indicate when the history was truncated. By default the current time and date is used and the history is fully deleted, but you can also use the optional UNTIL keyword with a timestamp value in the past which deletes data up to the specified time.

Any changes made to configuration values at the level of the file system – by scripts or tools such as `hdbnsutil` are also tracked. This is done at system startup when the system recognizes that offline configuration changes have been made and writes an entry into the log history; this entry does not include details of the values which have been changed.

Related Information

[Database-Specific Configuration Parameters \[page 293\]](#)

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

[SAP Note 2036111](#)

6.4.1 Database-Specific Configuration Parameters

In addition to the layers "default", "system", and "host", system configuration files also have a "database" layer to facilitate the configuration of properties for individual databases.

In general, you can configure database-specific properties both in the system database and in tenant databases themselves. Properties configured in the system database can be applied to all databases (if configured in the system layer) or to specific databases (if configured in database layer).

Properties configured in a tenant database apply to that tenant database only. Only properties in the following files can be configured in tenant databases:

- `attributes.ini`
- `docstore.ini`
- `dpserver.ini`
- `esserver.ini`
- `executor.ini`
- `extensions.ini`
- `global.ini`
- `indexserver.ini`
- `multidb.ini`
- `scriptserver.ini`
- `xsengine.ini`

File Location

If properties are configured in the database layer, a database-specific configuration file is stored at the following location on the server: `/hana/shared/$SID/global/hdb/custom/config/DB_<dbname>`

Example

The properties in the `nameserver.ini` file are not database specific. They can only be configured at system level. The `nameserver.ini` file is therefore stored at `/hana/shared/$SID/global/hdb/custom/config`.

However, the properties in the `indexserver.ini` can be database specific. Properties that are configured in the system layer and apply to all databases are stored in the `indexserver.ini` at `/hana/shared/$SID/global/hdb/custom/config`. Properties configured for an individual database override the

system-layer value and are stored in the `indexserver.ini` at `/hana/shared/$SID/global/hdb/custom/config/DB_<dbname>`.

Layered Configuration

Many properties can be configured in the system, host, and database layer. Values configured in the database layer take precedence over system-layer values.

However, when you are connected to a tenant database, you will see the database-layer value of a property is also displayed as the system-layer value. This is because from the perspective of the tenant database, the database and the system are effectively the same. The true system-layer value (that is, the value configured for all databases in the system database) is displayed in the tenant database as the default-layer value.

Values configured in the host layer take precedence over database-layer values. Host values can only be configured in the system database.

The following figure illustrates how layered configuration work. See also the example below.

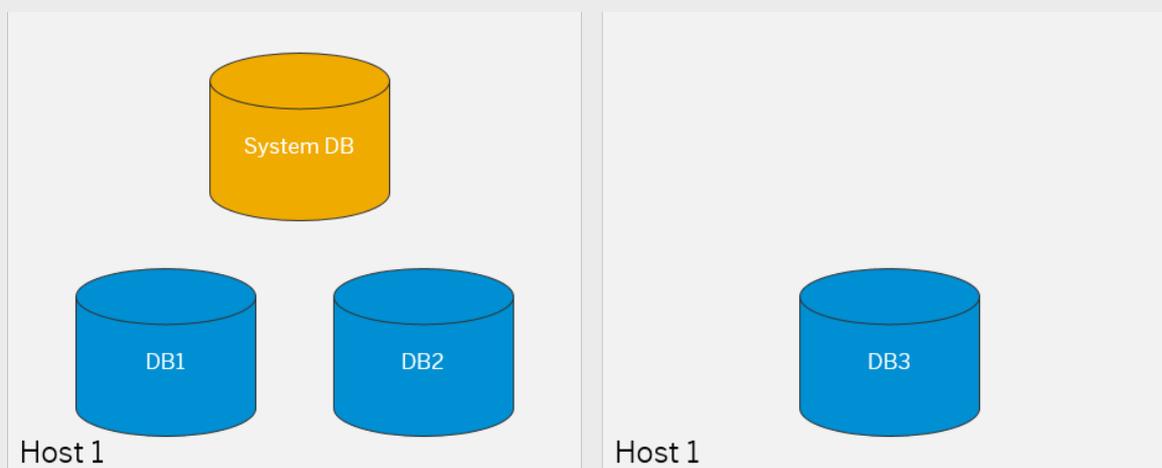


To view actual configuration values, open *Configuration of System Properties* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview* or query the following system views:

- M_INIFILE_CONTENTS (SYS_DATABASES)
This view can be accessed only from the system database. It contains the values configured for all properties on system, host, and database layer for **all** active databases.
- M_INIFILE_CONTENTS (SYS)
This view is available in every database and contains the values that apply to the database in question. Values that were configured in the system layer in the system database are identified as default-layer values. Values that were configured in the database layer in the tenant database are identified as system- and database-layer values. Values configured at the host layer are shown only for hosts on which the database is running.

❁ Example

A system has 3 tenant databases DB1, DB2, and DB3, distributed across 2 hosts Host A and Host B:



The default value of the property `[execution] max_concurrency` in the `global.ini` file is 0. The system administrator changes the default configuration of this property in the `indexserver.ini` file as follows:

First, the system administrator creates a new system-layer value (**10**) in `indexserver.ini`. Since the system-layer value applies to all tenant databases and cannot be changed by a tenant database user, users on all tenant databases initially see the value 10 as the default configuration:

Executed in DB1, DB2, and DB3:

```
SELECT * FROM SYS.M_INIFILE_CONTENTS WHERE KEY='max_concurrency'
```

FILE_NAME	LAYER_NAME	TENANT_NAME	HOST	SECTION	KEY	VALUE
global.ini	DEFAULT			execution	max_concurrency	0
indexserver.ini	DEFAULT			execution	max_concurrency	10 ← Effective value

Next, the system administrator sets a new value (**20**) for DB1, while leaving the configuration for DB2 and DB3 unchanged.

Executed in DB1:

```
SELECT * FROM SYS.M_INIFILE_CONTENTS WHERE KEY='max_concurrency'
```

FILE_NAME	LAYER_NAME	TENANT_NAME	HOST	SECTION	KEY	VALUE
global.ini	DEFAULT			execution	max_concurrency	0
indexserver.ini	DEFAULT			execution	max_concurrency	10
indexserver.ini	SYSTEM			execution	max_concurrency	20
indexserver.ini	DATABASE			execution	max_concurrency	20

← Effective value

Executed in DB2 and DB3:

```
SELECT * FROM SYS.M_INIFILE_CONTENTS WHERE KEY='max_concurrency'
```

FILE_NAME	LAYER_NAME	TENANT_NAME	HOST	SECTION	KEY	VALUE
global.ini	DEFAULT			execution	max_concurrency	0
indexserver.ini	DEFAULT			execution	max_concurrency	10

← Effective value

Note

In DB1, the database-layer value is duplicated to the system layer because from the perspective of the tenant database, the database and the system are effectively the same.

Finally, the system administrator sets a new value (15) for host A. Since host values take precedence over database values, this changes the effective value for DB1 and DB2.

Executed in DB1:

```
SELECT * FROM SYS.M_INIFILE_CONTENTS WHERE KEY='max_concurrency'
```

FILE_NAME	LAYER_NAME	TENANT_NAME	HOST	SECTION	KEY	VALUE
global.ini	DEFAULT			execution	max_concurrency	0
indexserver.ini	DEFAULT			execution	max_concurrency	10
indexserver.ini	HOST		HOST A	execution	max_concurrency	15
indexserver.ini	SYSTEM			execution	max_concurrency	20
indexserver.ini	DATABASE			execution	max_concurrency	20

← Effective value

Executed in DB2:

```
SELECT * FROM SYS.M_INIFILE_CONTENTS WHERE KEY='max_concurrency'
```

FILE_NAME	LAYER_NAME	TENANT_NAME	HOST	SECTION	KEY	VALUE
global.ini	DEFAULT			execution	max_concurrency	0
indexserver.ini	DEFAULT			execution	max_concurrency	10
indexserver.ini	HOST		HOST A	execution	max_concurrency	15
indexserver.ini	SYSTEM			execution	max_concurrency	20

← Effective value

Executed in DB3:

```
SELECT * FROM SYS.M_INIFILE_CONTENTS WHERE KEY='max_concurrency'
```

FILE_NAME	LAYER_NAME	TENANT_NAME	HOST	SECTION	KEY	VALUE
global.ini	DEFAULT			execution	max_concurrency	0
indexserver.ini	DEFAULT			execution	max_concurrency	10

← Effective value

Related Information

[SAP Note 2036111](#)

6.4.2 Configuring System Properties in SAP HANA Cockpit

From the cockpit overview of a specific system, you can drill down to view configuration (*.ini) files for the system, and add sections to the files.

Open *Configuration of System Properties* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview*.

Configuration files are separated into sections; sections bundle properties of the same category. Properties can be configured at different levels or layers depending on the configuration file. The following layers are available:

Layer	Description
Default	The default value for the property
System	The system-specific value for the property (configurable in the system database) If a system-specific value is not configured for a property, the default value applies.
Database	The database-specific value for the property (configurable in the system or tenant database) For some properties, it is possible to set database-specific values. If a database-specific value is not configured, the system-specific or host-specific value applies.
Host	The host-specific value for the property (configurable in the system database) For some properties, it is possible to set host-specific values for multiple-host systems. If a host-specific value is not configured for a property that can be set at host level, the system-specific value applies.

You can change configuration values (*Change Layer*) or quickly modify a value or assign a default.

You can use drop-down menus to select the *Configuration File* and the *Section* and *Host* in order to display specific configuration file contents.

In general, we do not recommend that you change the default values of parameters unless stated in the documentation or instructed to do so by SAP Support. For more information about frequently used parameters, see SAP Note 2036111.

While most parameters can be changed when the database is running, changes to some parameters require a database restart to take effect. To find out whether or not a restart is required for frequently used parameters, refer to SAP Note 2036111.

Related Information

[Database-Specific Configuration Parameters \[page 293\]](#)

[SAP Note 2036111](#)

[Add a System Property Section in SAP HANA Cockpit \[page 298\]](#)

[Add a System Property Parameter in SAP HANA Cockpit \[page 298\]](#)

[Modify a System Property in SAP HANA Cockpit \[page 299\]](#)

[Restore a System Property Default in SAP HANA Cockpit \[page 300\]](#)

6.4.2.1 Add a System Property Section in SAP HANA Cockpit

In the configuration files of an SAP HANA system, sections bundle properties of the same category.

Prerequisites

- Your database user has the system privilege INIFILE ADMIN.
- You must have the system privilege DATABASE ADMIN to change the system property of a tenant database from the system database.

Procedure

1. Open *Configuration of System Properties* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview*.
2. (Optional) Use drop-down menus to select the *Configuration File* in order to display specific configuration file contents.
3. Select *Add Section*.
4. Select a file from the drop down list.
5. Enter the name of the section.
6. From the drop down list, select the layer to *Add Values To*.
7. Enter a key and a value, and select *Add New Pair*.
8. Enter additional key value pairs as required.
9. Click *OK*.

6.4.2.2 Add a System Property Parameter in SAP HANA Cockpit

In the configuration files of an SAP HANA system, you can add a parameter.

Prerequisites

- Your database user has the system privilege INIFILE ADMIN.
- You must have the system privilege DATABASE ADMIN to change the system property of a tenant database from the system database.

Procedure

1. Open *Configuration of System Properties* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview*.
2. Use drop-down menus to select the *Configuration File* in order to display specific configuration file contents.
3. Select plus icon to add a parameter.
4. From the drop down list, select the layer to *Add Values To*.
5. Enter a key and a value, and select *Add New Pair*.
6. Enter additional key value pairs as required.
7. Click *OK*.

6.4.2.3 Modify a System Property in SAP HANA Cockpit

In the configuration files of an SAP HANA system, you can modify parameter values.

Prerequisites

- Your database user has the system privilege INIFILE ADMIN.
- You must have the system privilege DATABASE ADMIN to change the system property of a tenant database from the system database.

Context

i Note

In general, we do not recommend changing the default values of parameters unless stated in the documentation or instructed by SAP Support. For more information about frequently used parameters, see SAP Note 2036111.

Procedure

1. Open *Configuration of System Properties* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview*.
2. Use drop-down menus to select the *Configuration File* and the *Section* in order to display specific configuration file contents.
All the parameters in the section are listed. You can identify parameters that can have user-defined values by the edit (pencil) icon at the end of the row.

3. Select the edit icon for the value you wish to modify.
4. Enter the new value and click [Save](#).

Related Information

[SAP Note 203611](#)

6.4.2.4 Restore a System Property Default in SAP HANA Cockpit

In the configuration files of an SAP HANA system, you can reset changed parameters to their default values.

Prerequisites

- Your database user has the system privilege INIFILE ADMIN.
- You must have the system privilege DATABASE ADMIN to change the system property of a tenant database from the system database.

Procedure

1. Open [Configuration of System Properties](#) in SAP HANA cockpit by clicking the corresponding [Administration](#) link in the system [Overview](#).
2. Use drop-down menus to select the [Configuration File](#) and the [Section](#) in order to display specific configuration file contents.
All the parameters in the section are listed. You can identify parameters that have user-defined values by the edit (pencil) icon at the end of the row.
3. Select the edit icon for the user-defined value you wish to reset.
4. Select [Restore Default](#).

The user-defined value is cleared and the default value is re-applied.

6.4.3 Configuring System Properties in SAP HANA Studio

Use the SAP HANA studio to configure SAP HANA system properties (INI files). You can access and edit configuration files on the *Configuration* tab of the Administration editor.

Related Information

[Change a System Property in SAP HANA Studio \[page 301\]](#)

[Reset a System Property in SAP HANA Studio \[page 303\]](#)

6.4.3.1 Change a System Property in SAP HANA Studio

The properties of an SAP HANA system are defined in the parameters of its configuration files. Configuration files are separated into sections; sections bundle parameters of the same category.

Prerequisites

- You have the system privilege INIFILE ADMIN.
- In multiple-container systems, you must have the system privilege DATABASE ADMIN if you are changing the system property of a tenant database from the system database.

Procedure

1. In the Administration editor, choose the *Configuration* tab.
A list of all configuration files appears.
2. Expand the configuration file that you want to change.
All the sections of the configuration file are listed.
3. Expand the required section.
All the parameters of the section are listed. For each parameter, you can see the default value.
4. In the context menu of the configuration parameter that you want to change, choose *Change...*
The *Change Configuration Value* dialog box appears.
5. Enter the new value for the required layer.

Layer	Description
System	The value configured for the system applies to the system as whole, including all hosts of multi-host systems and all tenant databases of multi-DB systems.

Layer	Description
Host	<p>For some properties, it is possible to set host-specific values if the system has multiple hosts.</p> <p>If host-specific values are possible, you can expand the <i>Hosts</i> area of the <i>Change Configuration Value</i> dialog box, select the relevant host(s), and enter the host-specific value(s).</p> <p>It is possible to enter both a value for the system as a whole and for individual hosts. In this case, the system-specific value only applies to those hosts that do not have a host-specific value.</p>
Database	<p>For some properties, it is possible to set database-specific values if the system has tenant databases.</p> <p>If database-specific values are possible for a given property, they can be configured both in the system database and the tenant database.</p> <p>From the system database, you can configure database-specific values for all tenant databases in the system. From a tenant database, you can configure database-specific values only for that database.</p> <p>It is possible to enter a value for the system as a whole and individual databases. In this case, the system-specific value only applies to those databases that do not have a database-specific value.</p>

Note

If it is not possible to enter a host-specific or a database-specific value, the disabled icon () is displayed in the host or database column of the list view, and there is no *Hosts/Databases* area in the *Change Configuration Value* dialog box.

Results

- If you entered a new value for a parameter at system level, it is displayed in the *System* column with a green circle ().
- If you entered a new value for a parameter at host level, a gray rhombus () appears in the *Host* column. To show information about a specific host, select the host from the *Host* filter.
- If you entered a new value for a parameter at database level, a gray rhombus () appears in the *Database(s)* column. If you are logged on to the system database, you can show information about a specific database by selecting the database from the *Database* filter. This is only possible in the system database.

Next Steps

If necessary, restart the system.

6.4.3.2 Reset a System Property in SAP HANA Studio

You can restore changed parameters in the configuration files of an SAP HANA system back to their default values.

Prerequisites

- You have the system privilege INIFILE ADMIN.
- In multiple-container systems, you must have the system privilege DATABASE ADMIN if you are resetting the system property of a tenant database from the system database.

Procedure

1. In the Administration editor, choose the *Configuration* tab.
A list of all configuration files appears.
2. Expand the configuration file that you want to change.
All the sections of the configuration file are listed.
3. Expand the required section.
All the parameters in the section are listed. You can identify parameters that have user-defined values at system level and/or host level and/or database level with a green circle (●) and gray rhombus (◆) respectively.
4. To delete a user-defined value and restore the default value, you can choose one of the following methods:

Procedure	Result
<p>Delete with automatic reset:</p> <ol style="list-style-type: none">1. In the context menu of the configuration parameter, choose <i>Delete</i>. The <i>Delete Configuration Value</i> dialog box appears.2. Choose the layer whose user-defined values you want to delete.3. Choose <i>Delete</i>.	<p>The user-defined value(s) are cleared and the default value(s) are re-applied.</p> <div data-bbox="826 1541 1391 1675"><p>i Note If you added a new parameter to a section, <i>Delete</i> deletes the parameter.</p></div>
<p>Manually restore default:</p> <ol style="list-style-type: none">1. In the context menu of the configuration parameter, choose <i>Change...</i> The <i>Change Configuration Value</i> dialog box appears.2. For the required layers, choose <i>Restore Default</i>, or if you want to reset all visible layers, choose <i>Restore Default for All</i>.3. Choose <i>Save</i>.	<p>The user-defined value(s) are cleared and the default value(s) are re-applied.</p>

6.4.4 Configure System Usage Type

You can configure the usage type of an SAP HANA system (for example, production, development) during installation with the `system_usage` parameter, or later by changing the system properties. Clients such as the SAP HANA cockpit and SAP HANA studio can use this property to alter behavior.

Prerequisites

You have the system privilege INFILE ADMIN.

Context

The system usage type also governs which resource group (production, test, development) to which the system will be automatically assigned in SAP HANA cockpit when you register the system as a cockpit resource.

Procedure

1. In the `global.ini` configuration file, change the value of the `usage` property in the `system_information` section.

You can enter any value, but the values listed below can be used by clients to alter behavior. This may trigger a warning to users when they are about to perform critical operations on systems with usage type **production** (for example, execute SQL statements, stop or restart the system, perform a data backup, and so on).

- production
- test
- development
- custom (default)

i Note

If the system usage type is anything other than one of these values, you will not be able to register the SAP HANA system as a resource in the SAP HANA cockpit.

2. Restart the system.

Related Information

[Managing Resource Groups \[page 80\]](#)

6.4.5 Reserve Connections for Administrators

If the maximum number of connections has been reached in an SAP HANA system, it is not possible for anyone to log on, not even an administrator. For this reason, you can reserve a certain number of connections for administrative access only.

Prerequisites

You have the system privilege INFILE ADMIN and SESSION ADMIN.

Procedure

1. In the `global.ini` configuration file, change the value of the `reserved_connections` property in the `session` section.

As the value, specify the number of connections you want to reserve. The default number of reserved connections is 10. The minimum number is 1.
2. Restart the system.

Results

When the maximum number of connections minus the number reserved connections is reached, only an administrator with the system privilege SESSION ADMIN can log on to the system, for example, to resolve blocking situations by canceling sessions.

6.5 Managing SAP HANA Licenses

A valid license key is required to use the SAP HANA database. Additional license keys are required for certain applications running on SAP HANA, as well as and certain SAP HANA options and capabilities.

Related Information

[License Keys for the SAP HANA Database \[page 306\]](#)

[Managing Licenses in SAP HANA Cockpit \[page 308\]](#)

[Managing Licenses in SAP HANA Studio \[page 314\]](#)

6.5.1 License Keys for the SAP HANA Database

At least one license key is required to use the SAP HANA system. This license key must be installed in the system database. There are two kinds of license key: temporary license keys and permanent license keys.

Temporary License Keys

A temporary license key, which is valid for 90 days, is automatically installed in the system database of a new SAP HANA system and is effective for all tenant databases. During this period, you should request and install a permanent license key.

Permanent License Keys

You can request a permanent license key on SAP Support Portal (<http://support.sap.com>) under *Request a Key*. Permanent license keys are valid until the predefined expiration date. Furthermore, they specify the amount of memory licensed to the target SAP HANA database. Before a permanent license key expires, you should request and apply a new permanent license key. If a permanent license key expires in the system database, a temporary license key valid for 28 days is automatically installed. During this time, you can request and install a new permanent license key, for example, using the SAP HANA cockpit.

License Keys for Tenant Databases

You can install permanent license keys in individual tenant databases. The license key installed in a tenant database is valid for that database only and takes precedence over the license key installed in the system database. If a tenant-specific license key is not installed, the system database license key is effective in the tenant database.

→ Tip

The system view `SYS.M_LICENSE` provides tenant administrators with information on the license key effective in their tenant database, as well as where the license key is installed: in the tenant database itself or in the system database. System administrators can use the view `SYS_DATABASES.M_LICENSE` to see the same information for all tenant databases.

Unenforced and Enforced License Keys

There are two types of permanent license key available for SAP HANA: unenforced and enforced. If an unenforced license key is installed, the operation of SAP HANA is not affected if its memory consumption exceeds the licensed amount of memory. However, if an enforced license is installed, the database is locked down when the current memory consumption of SAP HANA exceeds the licensed amount of memory plus some tolerance. If this happens, either SAP HANA needs to be restarted, or a new license key that covers the amount of memory in use needs to be installed.

The two types of permanent license key differ from each other in the following line in the license key file:

License Key Type	License Key File Entry
Unenforced	SWPRODUCTNAME=SAP-HANA
Enforced	SWPRODUCTNAME=SAP-HANA-ENF
	SWPRODUCTNAME=SAP-HANA-DEV
	SWPRODUCTNAME=SAP-HANA-DIGITAL

Note

It is technically possible to install an enforced license in an SAP HANA database with a regular, unenforced permanent license. In this case, the unenforced license key has priority. That is, if a valid unenforced license key is found, excessive memory consumption will not result in a system lockdown. However, if one license key expires and becomes invalid, the other one, if valid, becomes the valid license key of the instance. If the latter is an enforced license key, then the memory consumption check is enforced.

Database Lockdown

The database goes into lockdown mode in the following situations:

System database

- The permanent license key has expired and either:
 - You did not renew the subsequently installed temporary license key within 28 days, or
 - You did renew the subsequently installed temporary license key but the hardware key has changed
- The installed license key is an enforced license key and the current memory consumption exceeds the licensed amount plus the tolerance.
- You deleted all license keys installed in the system database.

Tenant database

- The effective permanent license key has expired.
If the effective license key is installed in the system database, the conditions above apply.
- The effective license key is an enforced license key and the current memory consumption exceeds the licensed amount plus the tolerance.
If the effective enforced license key is installed in the tenant database, it takes precedence over the license key installed in the system database. The tenant database remains in lock-down mode, even if there is a valid license key available in the system database.

In lockdown mode, it is not possible to query the database. Only a user with the system privilege LICENSE ADMIN can connect to the database and execute license-related queries, such as, obtain previous license data, install a new license key, and delete installed license keys. The database cannot be backed up in lockdown mode.

i Note

In a locked-down tenant database, deleting the locally installed license key will resolve the situation, assuming the system database has a valid license and the locally installed license key is not an enforced license key.

Further SAP HANA Licenses

Additional licenses are required for certain applications running on the SAP HANA database, as well as certain features and capabilities of SAP HANA. For more information, see SAP Note 1644792.

Related Information

[SAP Support Portal](#) 

[Managing Licenses in SAP HANA Cockpit \[page 308\]](#)

[SAP Note 1644792](#) 

6.5.2 Managing Licenses in SAP HANA Cockpit

You can use the SAP HANA cockpit to see which licenses are available in your database, to install new license keys, and to view memory usage with respect to licensing.

Related Information

[View Licenses \[page 309\]](#)

[Install a Permanent License \[page 311\]](#)

[Export Usage Data \[page 312\]](#)

[Delete Licenses \[page 313\]](#)

6.5.2.1 View Licenses

You can view licenses installed in your SAP HANA database on the [License](#) page of the SAP HANA cockpit.

Context

You have the system privilege LICENSE ADMIN.

Procedure

On the [Overview](#) page, choose the [Manage system licenses](#) link.

Results

The [License](#) page opens. All licenses installed in the database are listed on the left. If you want to view the full details, including memory usage data, related to a particular license, simply click it. For more information, see [License Details](#).

i Note

If you are viewing license information in the system database, usage data is for the system as a whole. In a tenant database, only the usage data of the tenant is shown.

Related Information

[License Details \[page 310\]](#)

6.5.2.1.1 License Details

The [License](#) page provides you with detailed information about all licenses installed in the SAP HANA database.

General Information for SAP HANA Database Licenses

Field	Description
Hardware Key	Unique hardware key
System ID	Unique SAP system identifier
License Type	License type, either permanent or temporary
Product Description	SAP HANA database
Starts On	Date as of which the license is valid
Expires On	Date on which the license will expire
Licensed Memory Usage	Amount of memory usage licensed
Peak Memory Usage	Highest recorded value for main memory usage consumed by the SAP HANA database
GLAS ID	Unique product ID of SAP HANA required for license auditing

i Note

If you are viewing the information in the system database, memory usage data is for the system as a whole. In a tenant database, only the memory usage data of the tenant is shown.

The following additional fields are available if SAP Business Warehouse (SAP BW) is running on the SAP HANA database.

Field	Description
Peak Memory Usage of SAP BW	Highest recorded value for main memory usage consumed by SAP BW
GLAS ID of SAP BW	Unique product ID of SAP BW required for license auditing
Peak Memory Usage of Non-SAP BW Components	Highest recorded value for main memory usage consumed by SAP HANA
GLAS ID of Non-SAP BW Components	Unique product ID of SAP HANA required for license auditing

General Information for Other SAP HANA Licenses

Field	Description
Hardware Key	Unique hardware key

Field	Description
System ID	Unique SAP system identifier
License Type	License type, either permanent or temporary
Product Description	Product for which the license is valid
Starts On	Date as of which the license is valid
Expires On	Date on which the license will expire
Licensed <license_metric>	Amount of usage licensed The unit of measurement varies from product to product. For example, SAP HANA dynamic tiering licensing is based on hard disk usage.
Peak <license_metric>	Highest recorded usage value consumed
GLAS ID	Unique product ID required for license auditing

6.5.2.2 Install a Permanent License

To use SAP HANA, you must request and install a permanent license key. You can do this in the SAP HANA cockpit.

Prerequisites

- You have the necessary authorization to request permanent license keys on SAP Support Portal.
- To install a license key in the SAP HANA system, you have the system privilege LICENSE ADMIN.

Context

You need to request and install a permanent license key, for example, if the current license key of your SAP HANA database is about to expire, or you want to extend the amount of memory licensed for your database.

You can install license keys in the system database or in a tenant database. A license key installed in the system database is valid for any tenant database that doesn't have its own license.

i Note

You can also request and install the license keys required for applications running on SAP HANA, as well as SAP HANA options and capabilities, using the procedure described here.

Procedure

1. On the *Overview* page, choose the *Manage system licenses* link.

i Note

If the system is in lockdown, you will be automatically prompted to navigate to the *License* page.

2. Request a permanent license key as follows:

- a. On the *License* page, choose the *Request New License* button in the footer.

If the database is currently running on a temporary license key, the hardware key and the system ID are displayed. If the database already has a valid permanent license key, the installation number and system number are displayed. You will need this information to fill out the license key request form.

- b. Choose *Go to SAP Support Portal*.
- c. In SAP Support Portal, choose *Request Keys*.

You are forwarded to the license key application of the SAP ONE Support Launchpad, where you can request a new key. When completing the request form, if you have the installation number and system number, then enter them first so that the other input fields are auto-completed. When you have finished, choose *Submit*.

The permanent license will be sent to you as an e-mail attachment.

3. Install the license key by choosing *Upload New License* and uploading the license file (*.txt file) that you received by e-mail.

i Note

If you are installing a second or subsequent permanent license key, it must have the same system-identification data as the permanent license key previously installed in the database. In particular, the system ID (SID), hardware key, installation number, and system number must be the same. If any difference is detected in this data, the installation of the license key fails and no change is made to the license key in the database.

6.5.2.3 Export Usage Data

SAP HANA licensing is based on used memory. The SAP HANA database keeps track of the highest value for used memory per calendar month for a one-year period. You may be requested by SAP to export this data for license audit purposes to verify the suitability of your license.

Prerequisites

You have the system privilege LICENSE ADMIN.

Procedure

1. On the *Overview* page, choose the *Manage system licenses* link.
2. Click *Export System Measurement* in the footer toolbar.

Results

Usage data for all installed licenses is exported to the file `SAPHANASystemMeasurement.xml`, which is downloaded in line with your browser's file download settings.

i Note

In Safari, the file is not automatically downloaded. Instead, the content opens in a new tab or window and you must manually save the file by pressing `⌘` + `S`, choosing *page source*, and specifying a file name.

Related Information

[SAP Note 1704499](#) 

6.5.2.4 Delete Licenses

You can delete all existing license keys in the SAP HANA database, for example, if permanent license keys with an incorrect installation number or incorrect system number were installed.

Prerequisites

You have the system privilege LICENSE ADMIN.

Procedure

1. On the *Overview* page, choose the *Manage system licenses* link.
2. Click *Delete All Licenses* in the footer toolbar.

Results

All permanent license keys are deleted.

If you are in the system database, this results in the lockdown of the system database and any tenant databases that do not have their own tenant-specific license key. You must install a new, valid permanent license key is required to unlock the database(s).

If you are in a tenant database, the license keys installed in the system database now take effect.

6.5.3 Managing Licenses in SAP HANA Studio

A valid license key is required to use the SAP HANA database. Additional license keys are required for certain applications running on SAP HANA, as well as and certain SAP HANA options and capabilities

Related Information

[License Keys for the SAP HANA Database \[page 306\]](#)

[Check the Current License Key \[page 314\]](#)

[Install a Permanent License \[page 315\]](#)

[Delete an Existing Permanent License Key \[page 316\]](#)

6.5.3.1 Check the Current License Key

You can check the properties of your SAP HANA license in the SAP HANA studio.

Prerequisites

You have the system privilege LICENSE ADMIN.

Procedure

In the *Systems* view, right-click the system and choose  *Properties*  *License* .

Results

On the *System License* page under *Current License Key*, the following information is available:

- License type
- Start date of the license key
- Expiration date of the license key

The *All Licenses* page provides information about any further licenses installed in the system, for example, for SAP HANA options.

Related Information

[Important Disclaimer for Features in SAP HANA Platform \[page 1980\]](#)

6.5.3.2 Install a Permanent License

To use the SAP HANA database, you must request and install a valid permanent license key. You can do this in the SAP HANA studio.

Prerequisites

You have the system privilege LICENSE ADMIN.

Procedure

1. Get the information required to request a permanent license key.

To request the first permanent license key for a newly installed SAP HANA system, you need to provide the hardware key and the system ID. To request a subsequent permanent license key, you need the installation number and system number of your SAP HANA system. You can get the required information in the SAP HANA studio as follows:

In the *Systems* view, right-click the system and choose ► *Properties* ► *License* ►.

If the system is currently running on a temporary license key, the *Request License Key* screen area displays the hardware key and the system ID. If the system already has a valid permanent license key, the installation number and system number are displayed. Alternatively, you can use SQL to access the required information from the M_LICENSE system view.

2. Request a license key on SAP Support Portal (<http://support.sap.com>) by choosing *Request Keys*.

When completing the request form, if you have the installation number and system number, then enter them first so that the other input fields are auto-completed. When you have finished, choose *Submit*.

Permanent licenses are sent as e-mail attachments.

3. Install the license key
 - a. In the *Systems* view, right-click the system and choose ► *Properties* > *License* ⌵.
 - b. In the *Request License Key* area of the *System License* page, choose *Install License Key* and select the file that you received by e-mail.

i Note

If you are installing a second or subsequent permanent license key, it must have the same system-identification data as the permanent license key previously installed in the database. In particular, the system ID (SID), hardware key, installation number, and system number must be the same. If any difference is detected in this data, the installation of the license key fails and no change is made to the license key in the database.

→ Tip

You can also install the license key by executing the SQL statement `SET SYSTEM LICENSE '<license file content>'`. If you execute this statement use HDBSQL, you need to enable multiple-line mode before executing the statement.

Related Information

[SAP Support Portal](#) 

[SAP HANA HDBSQL \(Command-Line Reference\) \[page 1963\]](#)

[Run Long Commands in Multiple-Line Mode \[page 1975\]](#)

6.5.3.3 Delete an Existing Permanent License Key

Using the SAP HANA studio, you can delete all existing license keys in an SAP HANA database, for example, if permanent license keys with an incorrect installation number or incorrect system number were installed on the database.

Prerequisites

You have the system privilege LICENSE ADMIN.

Procedure

1. In the *Systems* view, right-click the system and choose ► *Properties* > *License* ⌵.

2. On the *System License* page choose *Delete License Key*.

→ Tip

You can also delete license keys by executing the SQL statement `UNSET SYSTEM LICENSE ALL`.

Results

All permanent license keys are deleted. This results in the lockdown of the database. The installation of a new, valid permanent license key is required to unlock the database.

Related Information

[Execute SQL Statements in SAP HANA Studio \[page 118\]](#)

[SAP HANA HDBSQL \(Command-Line Reference\) \[page 1963\]](#)

6.6 Monitoring the SAP HANA Database

It's important that you monitor the operation of SAP HANA databases on a regular basis. Although SAP HANA actively alerts you of critical situations, keeping an eye on resource usage and performance will help you identify patterns, forecast requirements, and recognize when something is wrong. You can monitor SAP HANA using both the SAP HANA cockpit and the SAP HANA studio. These tools rely on the monitoring and alerting information provided by system views and the statistics service.

Related Information

[Monitoring in SAP HANA Studio \[page 358\]](#)

[Monitoring in SAP HANA Cockpit \[page 317\]](#)

[System and Statistics Views \[page 388\]](#)

[The Statistics Service \[page 389\]](#)

6.6.1 Monitoring in SAP HANA Cockpit

You can perform several database monitoring tasks in the SAP HANA cockpit using a range of dedicated apps.

- Monitor overall database health

- Monitor status and resource usage of individual database services
- Analyze database performance across a range of key performance indicators related to memory, disk, and CPU usage
- Analyze the comparative memory utilization of column tables
- Analyze the memory statistics of the components of database services
- Monitor alerts occurring in the database and analyze patterns of occurrence
- Configure the alerting mechanism, for example, change alert threshold values, switch alert checkers on/off, and check for alerts out of schedule
- Monitor the status of system replication (if enabled)

If your system contains tenant databases and you are the overall system administrator, you can perform additional monitoring tasks at the system level.

i Note

If your system does not support tenant databases (single-container system), system and database are perceived as a single unit and are monitored and administered as a single database.

Related Information

[Open SAP HANA Cockpit \[page 47\]](#)

[Managing Tenant Databases \[page 189\]](#)

[Monitoring and Analyzing with the Performance Monitor \[page 397\]](#)

[Monitoring Alerts \[page 341\]](#)

[Configuring Alerts \[page 346\]](#)

6.6.1.1 Using the Overview to Manage a Resource

The *Overview* displays important metrics and available functions, regardless of whether the resource you are managing is a single-host or multi-host system, or a tenant database.

Through the *Overview*, you can view key health indicators for this specific resource, such as database status, alerts, and resource utilization. You also have access to tools that allow you to perform database administrations tasks, such as performance analysis, and executing SQL statements. Different parts of a single tile can link to different views or applications. This way, you can see various components in a single view and make the decision whether to further examine issues by drilling down.

To launch the *Overview*, drill down on the name of the resource from *My Resources* or from a group. Unless your administrator has enabled single sign-on, you'll need to connect to the resource with a database user that has the system privilege CATALOG READ and SELECT on `_SYS_STATISTICS`.

i Note

If you register an offline resource and display its *Overview*, you'll notice that much of the information that appears on the *Overview* of an online resource is missing. You can use the cockpit to start the resource: click the *Stopped* icon, then click *Start System* (upper right).

Overall Database Status

Overall Database Status can be running, running with issues, or stopped. Clicking on this status brings you to *Manage Services* where you can stop or kill a service, and start or stop a system.

Alerts

Alert counts for the resource are displayed for high- and medium-priority alerts, broken down by the nine alert categories defined in SAP HANA. (You can refresh the displayed data by using the manual or auto-refresh icons in the top right corner). Clicking on an *Alerts* tile brings you to the *Alert Monitor* for the resource. In the bottom right corner there is a status message showing vital information about SAP HANA processes that collect data. By noting the status messages within the tile, you can easily ascertain the validity of what you are seeing.

Usage and Performance Metrics

You can monitor key database metrics through the *CPU Usage*, *Memory Usage* and *Disk Usage* tiles, as well as the *Threads*, *Sessions* and *Monitor Statements* tiles. In a multi-host system, each host is represented by a clickable bar, with the selected host having a time graph displayed to the right of the bar chart. Hover over the bars to see details for the selected host. If a bar is highlighted, there is an associated high (red) or medium (yellow) alert. With single-host resources, since there is only one host, no bar graphs are displayed. By viewing this high-level information, you can decide whether to drill down to the *Performance Monitor*. See *Monitoring and Analyzing Past Performance*.

Smart Data Access

Without copying data directly into an SAP HANA database, you can use Smart Data Access to access remote data as if it were stored in local tables. Refer to the Smart Data Access section of the *SAP HANA Administration Guide*.

System Replication

If a database resource is part of a system replication configuration, you can monitor the status of replication between the primary system and the secondary system(s). See *Monitor System Replication*.

Additional Functionality

You can launch additional functionality by selecting any of the links organized under the headings *Monitoring*, *DB Administration*, *User and Role Management*, *Alerting and Diagnostics*, *Other Administration*, *Application*

[Lifecycle Management](#), [Platform Lifecycle Management](#) and [Help](#). Specific links and related tasks are described in the subsequent topics of this guide.

Security

The [Data Storage Security](#), [Auditing](#), [Authentication](#) blocks and the [Security Related Links](#) help you to monitor many critical security settings. Additionally, you can perform administration tasks related to data and communication encryption, and audit logging. See [Monitoring Critical Security Settings](#).

Performance Management

Use [Analyze Workload](#), [Capture Workload](#) and [Replay Workload](#) to manage performance. See [Capturing and Replaying Workloads](#).

SAP HANA Options

Other cockpit features that allow you to manage additionally-installed contexts (for example, SAP HANA dynamic tiering) are only visible and available if the specific context has been installed.

Related Information

[Monitoring Tenant Databases in SAP HANA Cockpit \[page 244\]](#)

[Manage Services \[page 321\]](#)

[Monitoring Alerts \[page 341\]](#)

[Monitoring and Analyzing with the Performance Monitor \[page 397\]](#)

[Monitoring SAP HANA System Replication with the SAP HANA Cockpit \[page 1184\]](#)

[Capturing and Replaying Workloads \[page 410\]](#)

[Managing Multiple Resources in SAP HANA Cockpit \[page 163\]](#)

[Connect to a Resource using Database Credentials \[page 168\]](#)

6.6.1.1.1 System General Information

Accessing general information about the SAP HANA system, such as operational status and database version, can assist you to monitor your system.

In SAP HANA cockpit, you can access [General Information](#) by drilling down in the system [Overview](#). To do this, your database user needs the system privilege CATALOG READ.

Details include:

- Information such as operational status, system usage type, whether the system has multiple hosts, the number of hosts (if distributed), and database version
- The SAP HANA version history
- Information about the plug-ins that are installed
- The status of replication from your production system to a secondary system. This information is only available and applicable if you are operating a secondary instance of your database (for example, in a high availability scenario). If this is the case, then content from the primary or production instance of your database is replicated to the secondary instance.

6.6.1.2 Monitoring Database Health and Resource Usage

To identify problems early and avoid disruptions, you need to monitor your SAP HANA database continuously.

You can monitor the overall status and resource usage of the SAP HANA database at a glance on the homepage of the SAP HANA cockpit. Then, drill down for more detailed monitoring and analysis.

6.6.1.2.1 Manage Services

To monitor the health of your SAP HANA database in more detail, for example, to troubleshoot performance bottlenecks, you can analyze the status and resource usage of individual database services. If necessary, you can perform follow-up operations, such as starting missing services, stopping a service, or killing a service. You can also start or stop a system.

Context

You can use the cockpit to monitor and manage more than one resource, each running version SAP HANA 1.0 SPS 12, or later. Any resource running version SAP HANA 2.0 SPS 01, or later is set in multiple-container mode, by default. The cockpit can also monitor single-container systems running earlier versions of SAP HANA. When you drill down to the system [Overview](#), and subsequently to [Manage Services](#), the operations you have the option to perform depend on whether you have drilled down through the system database or the tenant.

Procedure

Open [Manage Services](#) in SAP HANA cockpit by clicking the [Overall Database Status](#) in the system [Overview](#). You see the status of all the services in the database. For each service, detailed information about its memory consumption is available. For more information, see [Service Details](#).

i Note

Not all columns are visible by default. You can configure which columns are visible by clicking the configuration button in the table toolbar. You can configure the sort order of the information by clicking the sort button.

Next Steps

- If there are any alerts in the system, you can open them by clicking [Go to Alerts](#).
- If you want to investigate the memory usage history of a particular service, click the mini chart in the [Memory](#) column to open [Memory Analysis](#) for the service in a new window. See [Analyze Memory Allocation Statistics](#). You can also use [Reset Memory Statistics](#) link to clear the statistics history.
- Depending on the situation, you may need to perform further operations on all or selected services (for example, start, stop, or kill a service). For more information about the available options, see [Operations on Services](#).
- If necessary, you can also start or stop a system. See [Start a Resource](#) and [Stop a Resource](#).

Related Information

[Open SAP HANA Cockpit \[page 47\]](#)

[Service Details \[page 246\]](#)

[Operations on Services \[page 325\]](#)

[Monitoring Alerts \[page 341\]](#)

[Analyze Memory Statistics \[page 327\]](#)

[Assign Roles to a Database User \[page 799\]](#)

[Start a Resource \[page 179\]](#)

[Stop a Resource \[page 180\]](#)

[Add or Remove Services in a Tenant Database \[page 255\]](#)

6.6.1.2.1.1 Service Details

[Manage Services](#) provides you with detailed information about database services for an individual resource.

i Note

Not all of the columns listed below are visible by default. You can add and remove columns in the table personalization dialog, which you open by clicking the personalization icon in the table toolbar.

The table below lists the information available for services.

Column	Description
Host	Name of the host on which the service is running
Status	<p>The status of the service</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> • <i>Running</i> • <i>Running with Issues</i> (where at least one service is not running, or there is at least one high alert) • <i>Starting</i> • <i>Stopping</i> • <i>Not Running</i> <p>To investigate why the service is not running, you can navigate to the crash dump file created when the service stopped.</p> <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>The crash dump file opens in the <i>Trace</i> tool of the SAP HANA Web-based Development Workbench. For this, you need the role <code>sap.hana.xs.ide.roles::TraceViewer</code> or the parent role <code>sap.hana.xs.ide.roles::Developer</code>.</p> </div>
Service	Service name, for example, indexserver, nameserver, xsengine, and so on
Role	<p>Role of the service in a failover situation</p> <p>Automatic failover takes place when the service or the host on which the service is running fails.</p> <p>The following values are possible:</p> <ul style="list-style-type: none"> • <i>Master</i> The service is the active master worker. • No entry The service is a slave worker. • <i>Standby</i> The service is in standby mode. It does not contain any data and does not receive any requests.
Port	Port that the system uses for internal communication between services
Start Time	<p>Time at which the service started</p> <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>The time is given in the timezone of the SAP HANA server.</p> </div>
CPU	<p>Mini chart visualizing the CPU usage of the service</p> <p>Clicking the mini chart opens the <i>Performance Monitor</i> for a more detailed breakdown of CPU usage.</p>

Column	Description
Memory	<p>Mini chart visualizing the memory usage of the service</p> <ul style="list-style-type: none"> • Dark green shows the service's used memory. • Light green shows the service's peak memory. • The grey stroke represents the effective allocation limit. • The light grey background represents the physical memory. <p>Clicking the mini chart opens the Memory Analysis app for a more detailed breakdown of memory usage.</p>
Used Memory (MB)	<p>Amount of memory currently used by the service</p> <p>Clicking the mini chart opens the Memory Analysis app for a more detailed breakdown of memory usage.</p>
Peak Memory (MB)	Highest amount of memory ever used by the service
Effective Allocation Limit (MB)	Effective maximum memory pool size that is available to the process considering the current memory pool sizes of other processes
Memory Physical on Host (MB)	Total memory available on the host
All Process Memory on Host (MB)	Total used physical memory and swap memory on the host
Allocated Heap Memory (MB)	Heap part of the allocated memory pool
Allocated Shared Memory (MB)	Shared memory part of the allocated memory pool
Allocation Limit (MB)	Maximum size of allocated memory pool
CPU Process (%)	CPU usage of process
CPU Host (%)	CPU usage on host
Memory Virtual on Host (MB)	Virtual memory size on the host
Process Physical Memory (MB)	Process physical memory used
Process Virtual Memory (MB)	Process virtual memory
Shrinkable Size of Caches (MB)	Memory that can actually be freed in the event of a memory shortage
Size of Caches (MB)	Part of the allocated memory pool that can potentially be freed in the event of a memory shortage
Size of Shared Libraries (MB)	Code size, including shared libraries
Size of Thread Stacks (MB)	Size of service thread call stacks
Used Heap Memory (MB)	Process heap memory used
Used Shared Memory (MB)	Process shared memory used
SQL Port	SQL port number
Process ID	Process ID

Related Information

[Memory Usage in the SAP HANA Database \[page 470\]](#)

[Analyze Memory Statistics \[page 327\]](#)

[Monitoring and Analyzing with the Performance Monitor \[page 397\]](#)

6.6.1.2.1.2 Operations on Services

As an administrator, you may need to perform certain operations on all or selected services, for example, start missing services, or stop or kill a service.

You can perform several operations on database services from [Manage Services](#). You can trigger these operations by selecting the service and then clicking the required option in the footer toolbar.

i Note

To perform operations on services, you need to be an administrator. Depending on the service, some options may not be available. You can use the cockpit to monitor and manage more than one resource, each running version SAP HANA 1.0 SPS 12, or later. Any resource running version SAP HANA 2.0 SPS 01, or later is set in multiple-container mode, by default. The cockpit can also monitor single-container systems running earlier versions of SAP HANA. When you drill down to the system [Overview](#), and subsequently to [Manage Services](#), the operations you have the option to perform depend on whether you have drilled down through the system database or the tenant.

Option	Description
Start Missing Services	Starts any inactive services. Can only be performed on a tenant database if you drill down to Manage Services through the system database.
Stop Service	Stops the selected service normally The service is then typically restarted.
Kill Service	Stops the selected service immediately and if the related option is selected, creates a crash dump file The services is then typically restarted.
Add Service	Adds the service you select from the list. Can only be performed on a tenant database if you drill down to Manage Services through the system database. Services cannot be added to the system database itself.

i Note

To add a service, you must have the EXECUTE privilege on the stored procedure SYS. UPDATE_LANDSCAPE_CONFIGURATION.

Option	Description
Remove Service	<p>Removes the selected service. Can only be performed on a tenant database if you drill down to Manage Services through the system database.</p> <p>You can only remove services that have their own persistence. If data is still stored in the service's persistence, it is re-distributed to other services.</p> <p>You cannot remove the following services:</p> <ul style="list-style-type: none"> • Name server • Master index server • Primary index server on a host
Reset Memory Statistics	<p>Resets all memory statistics for all services. Can only be performed on a tenant database if you drill down to Manage Services through the system database.</p> <p>Peak used memory is the highest recorded value for used memory since the last time the memory statistics were reset. This value is useful for understanding the behavior of used memory over time and under peak loads. Resetting peak used memory allows you, for example, to establish the impact of a certain workload on memory usage. If you reset peak used memory and run the workload, then you can then examine the new peak used memory value.</p>
Go To Alerts	Displays the alerts for this database.

i Note

To remove a service, you must have the EXECUTE privilege on the stored procedure SYS. UPDATE_LANDSCAPE_CONFIGURATION.

i Note

The SAP HANA database provides several features in support of high availability, one of which is service auto-restart. In the event of a failure or an intentional intervention by an administrator that disables one of the SAP HANA services, the service auto-restart function automatically detects the failure and restarts the stopped service process. For more information about high availability, see *High Availability for SAP HANA* in the *SAP HANA Administration Guide*.

Related Information

[Memory Usage in the SAP HANA Database \[page 470\]](#)

[High Availability for SAP HANA \[page 1080\]](#)

[Start a Resource \[page 179\]](#)

[Stop a Resource \[page 180\]](#)

[Work with Alerts \[page 342\]](#)

6.6.1.2.2 Analyze Memory Statistics

Analyzing the memory allocation of the SAP HANA database can help you understand and resolve unusual memory usage and out-of-memory incidents.

Context

The [Memory Analysis](#) app enables you to visualize and explore the memory allocation of every service of a selected host during a specified time range. If you notice an increase in overall memory usage, you can investigate whether it's due to a particular component, allocator, or table.

Procedure

1. In SAP HANA cockpit, open [Memory Analysis](#) by selecting the corresponding link in the system [Overview](#).
2. Use the elements in the header to configure the display of memory statistics:

Element	Description
Host and Services Selection drop-down	Opt to display memory statistics from a different host and service combination.
Units Selection drop-down	Select a unit of measure in which to display all memory statistics
Time Selection	Select a time range for the memory allocation (upper) section, or select Custom in order to edit the time range.
Open In	Opt to view the service you have selected through the Memory Analysis app in the Performance Monitor or the Workload Analyzer.
Collection Frequency link	View how often information is collected to populate the various charts and graphs. If it is necessary, you can use the Configuration of System Properties app to modify the interval parameter of a statisticsserver component. Use caution because changes to the interval may affect performance.
Memory Alert Settings	Access these settings under the gear icon to turn on or off the display of alerts and specify which priority level of alerts should be displayed.
Threshold of Memory Usage	View the threshold of high priority memory usage alerts so you can determine what percentage of the effective allocation limit is being used by a service.

Note

Data is collected and displayed only after the statistics server has been enabled. Ensure that you have configured the system properties (the .ini file) so that the statistics server's `active` parameter been set to `true`.

- Analyze the memory statistics by exploring the data in the upper chart.

Element	Description
Allocated Memory	The pool of memory preallocated by the host for storing in-memory table data, thread stacks, temporary results and other system data structures.
Host Allocation Limit	The <code>global_allocation_limit</code> for the host (as set in the <code>global.ini</code> configuration file).
Service Allocation Limit	Each service running on the host has an allocation limit. Collectively, all services cannot consume more memory than the global allocation limit.
Total Used Memory	The total amount of memory used by the selected service on the selected host, including program code and stack, all data and system tables, and the memory required for temporary computations.
Memory Alert Icon	Represents one or more alerts triggered by a memory event. You can click individual alert icons and scroll through details, or navigate to the Alerts app.

- Move the vertical selection bar in the upper chart to populate the data in the lower chart. The vertical selection bar snaps to the closest time for which there is collected data for the components. You can control what is displayed in the lower chart by selecting one of its top tabs, as described in the next steps.
- Select the [Top Consumers](#) tab to have the lower chart display details about what is consuming the most memory for the host and service, in the given time period. Up to 50 top consumers are displayed. You can step through the collected data points by using the arrow buttons.

Element	Description
Consumer	Up to 50 allocators, objects or other elements that have consumed the most memory during the specific time (chosen by the vertical selection bar in the upper chart).
Used Memory	The amount of memory consumed, displayed in the measurement units specified at the top of the page.
Component	The component with which the consumer is associated.
Used memory history	Expand Used memory history to see a chart of the top consumers (up to 10) for the time range selected at the top of the page.

- Select the [Components](#) tab to have the lower chart display the [Used Memory by Component](#).

Element	Description
Used Memory by Component	For the specific time (chosen by the vertical selection bar in the upper chart), the components of the selected service are listed in descending order of used memory.
Used Memory by Type	The donut chart displays a visual representation of the types of used memory for the specific time.
Components Used Memory History	Filling the checkbox of one or more components populates the <i>Used Memory History</i> chart.

7. Select the *Allocators* tab to display more detailed memory use in the lower chart. You can filter by component type. You can step through the collected data points by using the arrow buttons. (Allocators that have used less than 1 GB of memory are not displayed.)

Element	Description
Used Memory by Allocator	For the specific time (chosen by the vertical selection bar in the upper chart), allocators of the selected component are listed in descending order of used inclusive memory. By clicking on a allocator, you can expand the list.
Filter by component name	To further refine the displayed allocator data, select the filter icon to specify one or more component name.
Allocators Used Memory History	Filling the checkbox of one or more allocators populates the Used Memory History chart.

8. Select the *Tables* tab to see statistics about the memory used by data tables in the lower chart.

Element	Description
Top Ten Tables by Size	Displays the breakdown of memory usage of the 10 highest consuming tables for the specific time (chosen by the vertical selection bar in the upper chart).
Top Ten Tables by Growth	Displays the memory usage of the 10 tables with the largest change in consumption for the selected time period. Hover over the data to see the <i>Previous Size</i> memory usage value from the beginning of the time period and the <i>Growth</i> during the time period (where the current size of the table is the sum of <i>Previous Size</i> and <i>Growth</i>).

9. Select the *Out of Memory Events* tab to have the lower chart display the number of unique out-of-memory events that have occurred in the time range specified in the header. (The vertical selection bar does not influence the number of events displayed.)

Element	Description
Occurrences	The number of times a specific OOM event has been triggered.
Last Occurrence	The time and date of the most recent occurrence of the OOM event.

Element	Description
Last Reason	The parameter that triggered the most recent occurrence of the OOM event.
Statement	The SQL statement related to the OOM event.
Statement Hash	The unique identifier for the OOM event. Click the OOM identifier to open the Workload Analyzer and investigate the event.

→ Tip

If an event has a corresponding OOM dump file, you can select [View Trace](#) to launch the Dump Viewer in the SAP Database Explorer.

In [Memory Statistics](#) charts you can choose to display historical data for a time range between 24 hours and six weeks. In order to have a date range longer than six weeks (42 days), you can use SQL to update the `RETENTION_DAYS_CURRENT` value in the table `"_SYS_STATISTICS"."STATISTICS_SCHEDULE"`.

Related Information

[Types of Memory Alerts \[page 330\]](#)

[System Views Used to Create Memory Alerts \[page 333\]](#)

[Memory Usage in the SAP HANA Database \[page 470\]](#)

[Manage Services \[page 321\]](#)

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

6.6.1.2.2.1 Types of Memory Alerts

The [Memory Statistics](#) chart can notify you about a variety of memory-related alerts.

Memory alerts related to the entire system

Alert	Description	Threshold Units	ALERT_ID
Licensed memory usage	Determines what percentage of licensed memory is used.	percent	44
Overflow of metadata version space	Determines the overflow ratio of the metadata version space.	ratio	74
Agent memory usage	Determines what percentage of total memory available on the agent is used .	percent	701

Memory alerts related to the host

Alert	Description	Threshold Units	ALERT_ID
Host physical memory usage	Determines what percentage of total physical memory available on the host is used. All processes that consume memory are considered, including non-SAP HANA processes.	percent	1
Streaming project physical memory usage	Determines what percentage of total physical memory available on the host is used for the streaming project.	percent	602
Plan cache hit ratio	Determines whether the plan cache hit ratio is too low.	ratio	91
Cached view size	Determines how much memory is occupied by the cached view.	percent	81

Memory alerts related to services

Alert	Description	Threshold Units	ALERT_ID
Memory usage of name server	Determines what percentage of allocated shared memory is being used by the name server on a host.	percent	12
Memory usage of services	Determines what percentage of its effective allocation limit is being used by a service.	percent	43
Plan cache hit ratio	Determines whether the plan cache hit ratio is too low.	ratio	91
Cached view size	Determines how much memory is occupied by the cached view.	percent	81

Memory alerts related to tables

Alert	Description	Threshold Units	ALERT_ID
Record count of non-partitioned column-store tables	Determines the number of records in non-partitioned column-store tables, but ignores information stored in statistics_exclude_tables. Current table size is not critical. Partitions need only be considered if tables are expected to grow rapidly. (A non-partitioned table cannot contain more than 2,000,000,000 (2 billion) rows).	records	17

Alert	Description	Threshold Units	ALERT_ID
Alert_Mon_Column_Tables_Record_Count_Incl	Determines the number of records in non-partitioned column-store statistics_exclude_tables. This alert is inactive, unless you activate it.	records	117
Table growth of non-partitioned column-store tables	Determines the growth rate of non-partitioned column tables.	percent	20
Record count of column-store table partitions	Determines the number of records in the partitions of column-store tables, but ignores information stored in statistics_exclude_tables. A table partition cannot contain more than 2,147,483,648 (2 billion) rows.	records	27
Alert_Partitioned_Table_Record_Count_Incl	Determines the number of records in the partitions of column-store statistics_exclude_tables. This alert is inactive, unless you activate it.	records	127
Size of delta storage of column-store tables	Determines the size of the delta storage of column tables.	MB	29
Total memory usage of column-store tables	Determines what percentage of the effective allocation limit is being consumed by individual column-store tables as a whole (that is, the cumulative size of all of a table's columns and internal structures).	percent	40
Memory usage of main storage of column-store tables	Determines what percentage of the effective allocation limit is being consumed by the main storage of individual column-store tables.	percent	45
Columnstore unloads	Determines how many columns in columnstore tables have been unloaded from memory. This can indicate performance issues.	tables	55

Alert	Description	Threshold Units	ALERT_ID
Total memory usage of table-based audit log	Determines what percentage of the effective memory allocation limit is being consumed by the database table used for table-based audit logging. If this table grows too large, the availability of the database could be impacted.	percent	64
Table growth of rowstore tables	Determines the growth rate of rowstore tables.	percent	67
Total memory usage of row store	Determines the current memory size of a row store used by a service.	percent	68
Row store fragmentation	Check for fragmentation of row store.		71
Overflow of rowstore version space	Determines the overflow ratio of the row store version space.	ratio	73
Rowstore version space skew	Determines whether the row store version chain is too long.	version	75
Auto merge of column-store tables	Determines whether the delta merge of a table was executed successfully.	records	88

6.6.1.2.2 System Views Used to Create Memory Alerts

The SAP HANA cockpit displays memory alerts in the *Memory Analysis* app based on SAP HANA system views.

The following system views provide the information with which values for current and historical memory allocation are calculated:

Views from the `_SYS` schema:

- `M_HOST_RESOURCE_UTILIZATION`
- `M_SERVICE_MEMORY`
- `M_SERVICE_COMPONENT_MEMORY`
- `M_RS_TABLES`
- `M_HEAP_MEMORY`
- `M_CS_COLUMNS`
- `M_OUT_OF_MEMORY_EVENTS`
- `M_SQL_PLAN_STATISTICS`

Views from the `_SYS_STATISTICS` schema:

- `HOST_RESOURCE_UTILIZATION_STATISTICS`

- HOST_SERVICE_MEMORY
- HOST_SERVICE_COMPONENT_MEMORY
- HOST_HEAP_ALLOCATORS
- GLOBAL_ROWSTORE_TABLES_SIZE_BASE
- HOST_COLUMN_TABLES_PART_SIZE
- STATISTICS_OBJECTS
- STATISTICS_CURRENT_ALERTS
- STATISTICS_ALERT_INFORMATION
- STATISTICS_ALERT_THRESHOLDS
- GLOBAL_OUT_OF_MEMORY_EVENTS

For more information about these views, see the *SAP HANA SQL and System Views Reference*.

6.6.1.2.3 Monitor Tables by Size and Usage

Monitor tables to optimize resource utilization and improve query performance.

Context

With [Table Usage](#) you can visualize tables by size, explore the usage history of tables, and move tables to warm storage.

Procedure

1. Open [Table Usage](#) in SAP HANA cockpit by clicking the corresponding [Monitor table usage](#) link in the system [Overview](#).
You see the status of the top column tables in the system by usage.
2. To filter tables shown, adjust the [Total Access/Size/Display](#) values and click [Go](#). Click [Reset](#) to remove filters.

For the best display, select up to 50 tables. Two options for table analysis are available:

- For table format display, choose  [table](#) ([Show table history as table](#)).
- For graphical format display, choose  [graph](#) ([Show table history as graph](#)). Mouse over a bubble to show usage per column table.

Next Steps

Monitor table operations to identify where you can improve performance and reduce memory utilization. Large in-memory tables that are accessed infrequently are good candidates for the SAP HANA dynamic tiering option. Note that tables moved into dynamic tiering disappear from table analysis displays.

Data Distribution Optimizer is part of the SAP HANA Data Warehousing Foundation option, which provides packaged tools for large scale SAP HANA use cases to support more efficient data management and distribution in an SAP HANA landscape. With Data Distribution Optimizer, SAP HANA Data Warehousing Foundation provides an SAP HANA XS-based tool to plan, adjust and analyze landscape redistribution. For more information, see *SAP HANA Data Warehousing Foundation - Data Distribution Optimizer Administration Guide* in Related Information.

Related Information

[Alert Details \[page 252\]](#)

[Alert Priorities \[page 253\]](#)

[Important Disclaimer for Features in SAP HANA Platform \[page 1980\]](#)

[SAP Note 2092669](#)

6.6.1.2.4 Configure Host Failover

For multi-host systems, you can configure host auto-failover so that if an active host fails, standby hosts take over to ensure the continued availability of the database.

Context

Host roles for failover are normally configured during installation. Using SAP HANA cockpit, you can monitor the status of individual hosts and switch the configured roles of hosts; you cannot increase or decrease the number of worker hosts and standby hosts in relation to each other.

The primary reason for changing the configured roles is to prepare for the removal of a host. In this case, change the configured role of the name server host to SLAVE and the configured role of the index server host to STANDBY before stopping the database instance on the host and removing the host.

i Note

To change host configuration, your database user must have the system privilege RESOURCE ADMIN and the object privilege EXECUTE on the procedure UPDATE_LANDSCAPE_CONFIGURATION.

Procedure

1. Open *Host Failover* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview*.
All the hosts in the system are displayed, whether or not they are operational, as well as additional information about their auto-failover status and configuration.

2. Click the gear button to customize which columns to display.

Column	Description
Host	Host name
Active	<p>Indicates the status of services running on the host.</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> ○ YES All services are active. ○ PARTIAL Some services are active. ○ STARTING Some services are active, some are starting. ○ STOPPING Some services are active, some are stopping. ○ NO No services are active.
Host Status	<p>Indicates the host's status and whether the system is operational.</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> ○ OK The system is operational and the host's actual role corresponds to its configured role. ○ IGNORE The system is operational. The host is configured as a standby host and is available, but not in use. ○ INFO The system is operational. The host's actual role is different from its configured role. ○ WARNING The system is not operational. The host will become available after start-up or failover. ○ ERROR The system is not operational. The host is missing.

Column	Description
Failover Status	<p>Displays the failover status so you can see which hosts are active and which are on standby.</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> ○ <Empty> Failover is neither active nor pending. ○ WAITING ... SEC The host has failed. The system is waiting to fail over. ○ WAITING The host has failed. The system is waiting for the host to restart to prevent unnecessary failover. ○ FAILOVER TO <host> The host has failed and failover to a target host is in progress. ○ FAILBACK TO <host> Failback to a worker host is in progress. This happens when the assigned standby host is stopped. However, there is no automatic failback while the standby host is still assigned since this would cause downtime. ○ FAILED Failover is not possible, for example, no further standby hosts available. For more information, see the nameserver trace.
Nameserver Role (Configured)	<p>Specifies the host's configured role as name server.</p> <p>The following roles are possible:</p> <ul style="list-style-type: none"> ○ MASTER 1, MASTER 2, MASTER 3 When you install a distributed system, up to three hosts are automatically configured as master name servers. The configured nameserver role of these hosts is MASTER 1, MASTER 2, and MASTER 3. ○ SLAVE Additional hosts in your system are configured as slave name servers. The configured nameserver role of these hosts is SLAVE.
Nameserver Role (Actual)	<p>Specifies the host's actual role as name server.</p> <p>The following roles are possible:</p> <ul style="list-style-type: none"> ○ MASTER During system start-up, one of the hosts configured as master name servers (that is, those hosts with the configured name server role MASTER 1, MASTER 2, or MASTER 3) is designated to be the active master name server. The actual nameserver role of this host is MASTER. This master name server assigns one volume to each starting index server (those with actual role MASTER or SLAVE), or no volume if it is a standby host (actual indexserver role STANDBY). If this active master nameserver host fails, one of the remaining hosts configured as a master nameserver becomes the active master name server. ○ SLAVE The actual nameserver role of the remaining hosts configured as master and slave hosts is SLAVE.

Column	Description
Indexserver Role (Configured)	<p>Specifies the host's configured role as index server.</p> <p>The following roles are possible:</p> <ul style="list-style-type: none"> ○ WORKER ○ STANDBY <p>When you install a distributed system, you can configure hosts either as WORKER or STANDBY index servers. A host configured as a standby index server is not used for database processing. All database processes run on the standby host, but they are idle and do not allow SQL connections.</p>
Indexserver Role (Actual)	<p>Specifies the host's actual role as index server.</p> <p>The following roles are possible:</p> <ul style="list-style-type: none"> ○ MASTER <ul style="list-style-type: none"> The actual master indexserver is assigned on the same host as the name server with the actual role MASTER. The actual index server role of this host is MASTER. The master index server provides metadata for the other active index servers (that is, those with actual indexserver role SLAVE). ○ SLAVE <ul style="list-style-type: none"> The actual index server role of remaining hosts (except those configured as standby hosts) is SLAVE. These are active index servers and are assigned to one volume. If an active index server fails, the active master name server assigns its volume to one of the standby hosts. ○ STANDBY <ul style="list-style-type: none"> The actual indexserver role of standby hosts is STANDBY. A standby host is not assigned a volume by the active master name server and it does not open an SQL port. <p>During normal operation when all hosts are available, a host with the configured role WORKER has the actual role MASTER or SLAVE, and a host with the configured role STANDBY has the actual role STANDBY. In the event of failover, the actual index server role of a host with the configured role STANDBY changes to SLAVE. The host status of the failed host changes from OK to INFO and the host status of the standby host changes from IGNORE to INFO.</p> <div style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>Failover is configured only for the name server and the index server on each host. The other components (for example, xsengine) are not configured individually as they are always failed over together with the index server.</p> </div>
Failover Group (Configured/Actual)	<p>A failover group can be defined for each host. In the event of failover, the name server tries to fail over to a host within the same group.</p>
Worker Groups (Configured/Actual)	<p>The worker groups (also referred to as host sub-roles) for the host can be set here. This is required to support heterogeneous hardware in the landscape which is required, for example, for the extension node feature.</p> <p>Worker groups may also be relevant in a single-host installation. The worker group name is a free text value that is validated to trap illegal characters.</p>

Column	Description
Host Roles (Configured)	<p>Specifies the host's configured database role.</p> <p>The following roles are possible:</p> <ul style="list-style-type: none"> ○ WORKER Worker host for database processing ○ STANDBY Standby host for database processing <p>Depending on your installation, the following additional host roles may be configured:</p> <ul style="list-style-type: none"> ○ EXTENDED_STORAGE_WORKER Worker host for SAP HANA dynamic tiering ○ EXTENDED_STORAGE_STANDBY Standby host for SAP HANA dynamic tiering ○ ETS_WORKER Worker host for SAP HANA accelerator for SAP ASE ○ ETS_STANDBY Standby host for SAP HANA accelerator for SAP ASE ○ RDSYNC Host for SAP HANA remote data sync ○ STREAMING Host for SAP HANA Streaming Analytics ○ XS_STANDBY Standby host for SAP HANA XS advanced runtime ○ XS_WORKER Host for SAP HANA XS advanced runtime
<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>Note</p> <p>Multiple host roles are not supported in production environments. However, if XS advanced runtime is installed, hosts can share multiple roles.</p> </div>	
Host Roles (Actual)	Specifies the host's actual database role
Storage Partition	<p>Specifies the number of the <code>mnt000...</code> subdirectory used by the host for storing data and logs, for example, <code>1</code> if the subdirectory is <code>mnt00001</code>, <code>2</code> if it is <code>mnt00002</code>, and so on.</p> <p>During installation, volumes for storing data and log files are defined. These are the directories where data and logs are stored. The default directories are:</p> <ul style="list-style-type: none"> ○ <code>/hana/data/<SID></code> for data ○ <code>/hana/log/<SID></code> for logs <p>Each active host has exactly one subdirectory beneath these directories called <code>mnt00001</code>, <code>mnt00002</code>, and so on. The next level in the file hierarchy is the actual volume, with one subdirectory for each service called <code>hdb00001</code>, <code>hdb00002</code>, and so on.</p> <p>In the event of failover, the volumes of the failed host are reassigned to the standby host.</p>

Column	Description
Removal Status	<p>Indicates the status of the table redistribution operation used to move data off the index server of a host that you plan to remove.</p> <p>Before you can remove an active host from a single-container system, you must move the tables on the index server of this host to the index servers on the remaining hosts in the system. Once the value in the <i>Removal Status</i> column changes to REORG FINISHED or REORG NOT REQUIRED, you can physically remove the host using the SAP HANA lifecycle management tool <code>hdblcm(gui)</code>.</p> <p>If your system is configured as a multiple-container system, you have to remove tenant-specific services first and then remove the host using the SAP HANA database lifecycle manager (HDBLCM). For more information, see <i>Remove a Service from a Tenant Database</i> in the <i>SAP HANA Administration Guide</i>.</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> ○ <Empty> Host has not been marked for removal. ○ REORG PENDING A redistribution operation is required to move tables to other hosts. ○ REORG ACTIVE A redistribution operation is in progress. For more information, you can query the system tables <code>SYS.REORG_OVERVIEW</code> and <code>SYS.REORG_STEPS</code>. ○ REORG FAILED A redistribution operation was executed and failed. For more information, query the system table <code>SYS.REORG_STEPS</code>. ○ REORG FINISHED A redistribution operation has completed. The host can be uninstalled. ○ REORG NOT REQUIRED A redistribution operation is not required. The host can be uninstalled.

- If you make changes to a configured role or configured group, click [Apply](#) so that your changes take effect.

6.6.1.3 Monitoring Health in Multi-Host Systems

For scale-out or multiple-host systems, SAP HANA cockpit provides status information on the health of system components on their respective servers, and on resource utilization of hardware components, including CPU, memory, network, and storage on the respective servers.

Open [System Health](#) in SAP HANA cockpit by clicking [Monitor system health](#) under [Monitoring](#) in the system [Overview](#) of a multi-host system.

The system health charts display the most recent 10 minutes of data, organized by hosts within the system. The host types include master, standby, dynamic tiering, streaming, and remote data sync. The health metrics are:

Metric	Details
Status	Status of the main service (indexserver)

Metric	Details
Critical alerts	Number of high priority alerts for the given host
CPU %	Percentage of CPU usage for the given host
Memory %	Percentage of memory usage for the given host
Out of Memory Events	Number of out of memory events
Unloads	Number of columns unloaded due to low memory
Disk Usage %	Percentage of disk usage for the given host
Network I/O	Number of network input/output events
Statements	Changes in number of statements per second
Versions	Number of multiversion concurrency control versions
Disk I/O	Number of disk input/output events

To reduce the number of hosts or show particular hosts side by side, click the  [settings](#) icon and select the desired hosts from the list.

By selecting specific system health information, you can drill down to details about alerts when critical alerts are present. You can also drill down to details on specific KPI, displayed in the Performance Monitor, and the Workload Analyzer.

Related Information

[Monitoring and Analyzing with the Performance Monitor \[page 397\]](#)

[Analyzing Workloads \[page 435\]](#)

6.6.1.4 Monitoring Alerts

As an administrator, you actively monitor the status of the system and its services and the consumption of system resources. However, you are also alerted of critical situations, for example: a disk is becoming full, CPU usage is reaching a critical level, or a server has stopped.

The internal monitoring infrastructure of the SAP HANA database is continuously collecting and evaluating information about status, performance, and resource usage from all components of the SAP HANA database. In addition, it performs regular checks on the data in system tables and views and when configurable threshold values are exceeded, issues alerts. In this way, you are warned of potential problems. The priority of the alert indicates the severity of the problem and depends on the nature of the check and configured threshold values. For example, if 90% of available disk space is used, a low priority alert is issued; if 98% is used, a high priority alert is issued. For more information about the technical implementation of monitoring and alerting features in SAP HANA, see *The Statistics Service*.

A summary of all alerts in the database is available on the homepage of the SAP HANA cockpit. To get more information about these alerts and to analyze the historical occurrence of alerts, you can drill down into the [Alerts](#) app.

In addition, several configuration options are available so that you can tailor alerting in the SAP HANA database to your needs (for example, changing alerting thresholds, switching particular alerts off, and setting up e-mail notification of alerts).

Related Information

[Configuring Alerts \[page 346\]](#)

[The Statistics Service \[page 389\]](#)

[Work with Alerts \[page 342\]](#)

[Search, Sort, and Filter Tools for Alerts \[page 343\]](#)

[Alert Details \[page 252\]](#)

[Alert Priorities \[page 253\]](#)

6.6.1.4.1 Work with Alerts

Find resources with alerts, drill down for information on active alerts, and use an alert checker.

Prerequisites

Your database user needs the object privilege SELECT on the schema `_SYS_STATISTICS`.

Procedure

1. Connect to the cockpit and sign in.

The URL takes this form:

```
https://<cockpit-host>:<port-number>
```

The port number was configured during cockpit installation.

2. Select a resource group at the top of the page.
3. Look at the Top Resources with Alerts tile. It displays the resources (up to five) that have the most high-priority alerts.
4. Click a resource on the Top Resources with Alerts tile to see all of that resource's alerts.
5. On the Alert Details page, click an alert in the left pane to display information.

If the list of alerts is long, use the [Search](#), Sort, Filter, and [Past Alerts](#) tools in the left pane to find particular alerts.

6. (Optional) Use the drop-down menu on the right to change the time period shown on the Past Occurrences of Alert graph.

7. (Optional) To display all past occurrences of the selected alert, click [Past Alerts](#) (bottom of left pane). To restore the list to current alerts only, click again.
8. (Optional) Click [Configure Alert](#) (lower right) to view, execute, or modify the alert checker for the alert currently displayed.
 - a. Review the details on the Alert Checker Configuration page.
 - b. Click [Edit](#) (lower right) to make changes.

For example, you can add e-mail addresses to receive alert notifications, or click the [Schedule Active](#) button to enable or disable the alert. Click [Save](#).
 - c. When you're satisfied with the settings, click [Check Now](#) (lower right).

The alert checker notifies you when the check is done. If the check generated alerts, you can click [Go to Alerts](#) to see them.

Related Information

[Search, Sort, and Filter tools for Alerts \[page 343\]](#)

[Open SAP HANA Cockpit \[page 47\]](#)

[Alert Details \[page 252\]](#)

[Alert Priorities \[page 253\]](#)

6.6.1.4.1.1 Search, Sort, and Filter Tools for Alerts

Use Search, Sort, and Filter tools to work with lists of alerts.

Search	Enter full or partial alert text in the Search box at the top of the screen and click the  search icon. The list is reduced to show only alerts that match your search string.
Change sorting rules	<p>You can modify the sorting rules to:</p> <ul style="list-style-type: none"> • List alerts in ascending (the default) or descending order • Sort by alert text, category, alert checker, priority, or date (the default). <p>Click the  sort icon, choose sorting rules, and click OK.</p>

Add a filter

You can filter the list of alerts by:

- Category - select as many as you want.
- Alert checker - select as many as you want.
- Priority - select as many as you want.
- Time frame - select one.
- Alerting hosts - select as many as you want.

Click the  *filter* icon. Select the type of filter (category, for example), then select one or more options or click *Select All*. Click *OK* to save the filter.

Modify a filter

Each active filter appears as a box in the blue bar at the top of the left pane. The box contains the filtering parameter

(priority or time frame, for example): .

To modify a filter, click its filtering parameter. The Filter By dialog opens.

Cancel a filter

To cancel a filter, click its  *cancel* icon in the blue bar at the top of the left pane.

6.6.1.4.2 Alert Details

When you select an alert, detailed information about the alert is displayed on the right.

The following detailed information about an alert is available:

Detail	Description
Category	The category of the alert checker that issued the alert Alert checkers are grouped into categories, for example, those related to memory usage those related to transaction management and so on.
Next Scheduled Run	When the related alert checker is next scheduled to run If the alert checker has been switched off (alert checker status <i>Switched Off</i>) or it failed the last time it ran (alert checker status <i>Failed</i>), this field is empty because the alert checker is no longer scheduled.
Interval	The frequency with which the related alert checker runs If the alert checker has been switched off (alert checker status <i>Switched Off</i>) or it failed the last time it ran (alert checker status <i>Failed</i>), this field is empty because the alert checker is no longer scheduled.

Detail	Description
Alerting Host & Port	Name and port of the host that issued the alert In a system replication scenario, alerts issued by secondary system hosts can be identified here. This allows you to ensure availability of secondary systems by addressing issues before an actual failover. For more information about monitoring secondary systems in SAP HANA, see <i>Monitoring Secondary Sites</i> in the <i>SAP HANA Administration Guide</i> .
Alert Checker	Name and description of the related alert checker
Proposed Solution	Possible ways of resolving the problem identified in the alert, with a link to the supporting app, if available
Past Occurrences of Alert	Configurable graphical display indicating how often the alert occurred in the past

Related Information

[Monitoring Secondary Sites \[page 1194\]](#)

6.6.1.4.3 Alert Priorities

The priority of an alert indicates the severity of the problem and how quickly action needs to be taken.

Priority	Description
Information	Action recommended to improve system performance or stability
Low	Medium-term action required to mitigate the risk of downtime
Medium	Short-term action required (few hours, days) to mitigate the risk of downtime
High	Immediate action required to mitigate the risk of downtime, data loss, or data corruption

6.6.1.4.4 Analyze Occurrences of an Alert Over Time

Analyzing when and how often an alert has occurred in the past can help you for example troubleshoot recurring problems and identify patterns.

Procedure

1. Open the *Alerts* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.

All latest alerts are displayed in list format on the left.

2. Find and select the alert that you want to analyze using the options available for filtering, searching, and sorting.
Detailed information about the alert is shown on the right, including a graph displaying how often the alert has been issued over time.
3. Select the timeframe that you want to analyze.
By default, the number of occurrences per hour over the last 24 hours are displayed.

Related Information

[Open SAP HANA Cockpit \[page 47\]](#)

[Assign Roles to a Database User \[page 799\]](#)

6.6.1.4.5 Configuring Alerts

Several configuration options are available so that you can tailor alerting in the SAP HANA database to your needs.

The following configuration options are possible:

- Change the threshold values that trigger alerts of different priorities.
- Set up e-mail notifications so that specific people are informed when alerts are issued.

In addition, you can perform the following actions on alert checkers:

- Run alert checkers on a once-off basis, regardless of their configured schedule or status
- Switch alert checkers off and on

Related Information

[Configure Alerting Thresholds \[page 348\]](#)

[Switch Alerting Off/On \[page 350\]](#)

[Set Up E-Mail Notification \[page 351\]](#)

[Check for Alerts Out of Schedule \[page 353\]](#)

[Alert Checker Details \[page 253\]](#)

[Alert Checker Statuses \[page 254\]](#)

6.6.1.4.5.1 Alert Checker Details

When you select an alert checker [Alert Configuration](#), detailed information about the alert checker and its configuration is displayed on the right.

The following detailed information about an alert checker is available:

Detail	Description
Header information	The name of the alert checker, its status, and the last time it ran
Description	Description of what the alert checker does, for example what performance indicator it measures or what setting it verifies
Alert Checker ID	The unique ID of the alert checker
Category	The category of the alert checker Alert checkers are grouped into categories, for example those related to memory usage, those related to transaction management, and so on.
Threshold Values for Prioritized Alerting	The values that trigger high, medium, low, and information alerts issued by the alert checker The threshold values and the unit depend on what the alert checker does. For example, alert checker 2 measures what percentage of disk space is currently used so its thresholds are percentage values.
	i Note Thresholds can be configured for any alert checker that measures variable values that should stay within certain ranges, for example, the percentage of physical memory used, or the age of the most recent data backup. Many alert checkers verify only whether a certain situation exists or not. Threshold values cannot be configured for these alert checkers. For example, alert checker 4 detects services restarts. If a service was restarted, an alert is issued.
Interval	The frequency with which the alert checker runs
Schedule Active	Indicator of whether the alert checker is running automatically at the configured interval
Proposed Solution	Possible ways of resolving the problem identified by the alert checker

Related Information

[Alert Checker Statuses \[page 254\]](#)

[Configure Alerting Thresholds \[page 348\]](#)

6.6.1.4.5.2 Alert Checker Statuses

The status of an alert checker indicates whether it is running on schedule, it has failed and been disabled by the system, or you switched it off.

Status	Description
Active	The alert checker is running on schedule.
Failed	<p>The alert checker failed the last time it ran (for example due to a shortage of system resources), so the system disabled it.</p> <p>The alert checker remains disabled for a specific length of time before it is automatically re-enabled. This length of time is calculated based on the values in the following columns of the table STATISTICS_SCHEDULE (_SYS_STATISTICS):</p> <ul style="list-style-type: none">• INTERVALLENGTH• SKIP_INTERVAL_ON_DISABLE <p>Once $INTERVALLENGTH \times SKIP_INTERVAL_ON_DISABLE$ has elapsed, the alert checker is re-enabled. The default values for all alert checkers are such that failed checkers remain disabled for 1 hour. The system determines the status of every alert checker and/or whether the time to re-enablement has elapsed every 60 seconds.</p> <p>You can also re-enable the alert checker manually by switching it back on in Alert Configuration.</p>
Switched Off	<p>You switched off the alert checker schedule.</p> <p>If you want the alert checker to run again automatically, you must manually switch it back on.</p>

Related Information

[Switch Alerting Off/On \[page 350\]](#)

[Configure Alerting Thresholds \[page 348\]](#)

6.6.1.4.5.3 Configure Alerting Thresholds

In many cases, you can configure the thresholds that trigger an alert. An alert checker can have a low, medium, and high priority threshold.

Context

Thresholds can be configured for any alert checker that measures variable values that should stay within certain ranges, for example, the percentage of physical memory used, or the age of the most recent data

backup. Many alert checkers verify only whether a certain situation exists or not. Threshold values **cannot** be configured for these alert checkers. For example, alert checker 4 detects services restarts. If a service was restarted, an alert is issued.

Procedure

1. Open *Alert Configuration* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview*.
2. Find the alert checker whose thresholds you want to change.
The detailed configuration of the alert checker is displayed on the right. For more information, see *Alert Checker Details*.
3. Open the alert checker for editing by clicking *Edit*.
4. Change the threshold values as required.
The threshold value depends on what the specific alert checker is measuring. For example, for alert checker 2 (disk usage), you could enter 90, 95 and 100 as the thresholds, where 90, 95, and 100 represent the percentage of disk space used.

→ Tip

The unit for the threshold value of the alert checker is indicated in brackets above the entry fields.

5. Save the alert checker.

Results

Alerts are issued when the alert checker records values that exceed the configured thresholds.

Related Information

[Open SAP HANA Cockpit \[page 47\]](#)

[Alert Checker Details \[page 253\]](#)

[Assign Roles to a Database User \[page 799\]](#)

6.6.1.4.5.4 Switch Alerting Off/On

If you no longer want a particular alert to be issued, you can switch off the underlying alert checker so it no longer runs automatically according to schedule. Alert checkers that the system has disabled must be switched back on manually.

Context

In some situations you may want to stop a particular alert from being issued, either because it is unnecessary (for example, alerts that notify you when there are other alerts in the system) or because it is not relevant in your system (for example, backup-related alerts in test systems where no backups are performed).

⚠ Caution

If you switch off alerts, you may not be warned about potentially critical situations in your system.

You can switch an alert checker back on again at any time.

You may also want to switch on alert checkers that the system has disabled, that is checkers with the status *Failed*. The system automatically disables alert checkers when they fail to run, for example, due to a shortage of system resources.

i Note

The system automatically switches failed alert checkers back on after a certain length of time. For more information, see *Alert Checker Statuses*.

It is possible to disable an alert for a particular table or schema. This is supported for the alerts "Record count of non-partitioned column-store tables" (ID 17) and "Table growth of non-partitioned column-store tables" (ID 20).

To exclude an alert to be issued for a particular table, use the following SQL statement:

```
INSERT INTO _sys_statistics.statistics_exclude_tables VALUES (<alert_id>,  
'<schema_name>', '<table_name>')
```

To exclude an alert to be issued for all tables of a particular schema, use the following SQL statement:

```
INSERT INTO _sys_statistics.statistics_exclude_tables VALUES (<alert_id>,  
'<schema_name>', null)
```

To re-enable the alerts, delete the entries from the table `_sys_statistics.statistics_exclude_tables`.

Procedure

1. Open *Alert Configuration* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview*.

2. Find the alert checker that you want to switch off or on.
The detailed configuration of the alert checker is displayed on the right. For more information, see *Alert Checker Details*.
3. Open the alert checker for editing by clicking *Edit*.
4. Set the *Schedule Active* switch control to *No* or *Yes*.
5. Save the alert checker.

Results

If you switched the alert checker off, its status changes to *Switched Off* and it is no longer scheduled to run automatically.

If you switched the alert checker on, its status changes to *Active* and it starts running again automatically according to its configured schedule.

Related Information

[Open SAP HANA Cockpit \[page 47\]](#)

[Alert Checker Details \[page 253\]](#)

[Alert Checker Statuses \[page 254\]](#)

6.6.1.4.5.5 Set Up E-Mail Notification

You can configure alert checkers so that you and other responsible administrators receive push notifications by e-mail when alerts are issued.

Context

If you want to be notified by e-mail about new alerts when they are issued, you can set this up in through *Alerts Configuration*. You can configure one or more default recipients to be notified when any alert checker issues an alert. In addition, if different people need to be notified about different alerts, you can configure dedicated recipients for these alert checkers.

Note the following behavior:

- If you configure checker-specific recipients, default recipient(s) will **not** be notified.
- If you delete all checker-specific recipients, default recipient(s) will be notified again, if configured.
- You can configure checker-specific recipients regardless of whether or not default recipients are configured.

Procedure

1. Open *Alert Configuration* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview*.
2. Configure the e-mail sender:
 - a. In the footer toolbar, choose *Configure Email*, then *Sender*.
 - b. Enter the following information for the e-mail sender:
 - Sender's e-mail address
E-mail address that is entered as the e-mail sender
 - SMTP server
The mail server that the system sends the e-mails to

i Note

The statistics service does not support a mail server that requires additional authentication.

- SMTP port
The default SMTP port is 25. If the configured mail server uses a different port, you must enter it.
3. Optional: Configure one or more default recipients.

Default recipients are notified about alerts generated by all alert checkers **except** those that have checker-specific recipients configured.

 - a. In the footer toolbar, click the envelope icon and choose *Default Recipient(s)*.
 - b. Enter the e-mail addresses of the recipients.
 - c. Save the configuration.
 4. Optional: Configure one or more recipients for specific alert checkers.

Checker-specific recipients are notified only about alerts generated by the alert checker in question. Default recipients (if configured) are not.

 - a. Find the alert checker that you want to configure.
The detailed configuration of the alert checker is displayed on the right.
 - b. Open the alert checker for editing by clicking *Edit*.
 - c. In the *Email* field, enter the e-mail addresses of the recipients.
 - d. Save the alert checker.

Results

The configured recipients will receive an email when an alert checker issues an alert. If the alert checker issues the same alert the next time it runs, no further e-mail is sent. However, when the alert checker runs and it does not issue an alert, indicating that the issue is resolved or no longer occurring, a final e-mail is sent.

Related Information

[Open SAP HANA Cockpit \[page 47\]](#)

[Alert Checker Details \[page 253\]](#)

6.6.1.4.5.6 Check for Alerts Out of Schedule

In general, alert checkers run automatically according to a configured schedule. If necessary, you can run an alert checker on a once-off basis outside of its schedule.

Context

In some cases, you may want to check for a particular alert outside of the alert checker's configured schedule. For example, to verify that the problem identified by a previous alert has been resolved.

Running an alert checker in this ad hoc way does not affect its configured schedule.

i Note

If you want to manually run an alert checker with the status *Switched Off* or *Failed*, you must switch it back on first.

Procedure

1. Open *Alert Configuration* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview*.
2. Find the alert checker that you want to run.
The detailed configuration of the alert checker is displayed on the right. For more information, see *Alert Checker Details*.
3. Choose *Check Now* in the footer toolbar.
The alert checker starts running. Once it has finished, you will be notified of the result.

Related Information

[Open SAP HANA Cockpit \[page 47\]](#)

[Alert Checker Details \[page 253\]](#)

[Alert Checker Statuses \[page 254\]](#)

[Switch Alerting Off/On \[page 350\]](#)

6.6.1.5 Monitor Disk Volume

In order to ensure that the database can always be restored to its most recent committed state, you can use the SAP HANA cockpit to check disk statistics to check that there is enough space on disk for data volumes and log volumes.

Context

A disk has multiple volumes. Each volume has a data volume and a log volume. Data volumes have one file (datavolume_0000.dat). Log volumes often have hundreds of files (multiple logsegment_000_00000003.dat; single logsegment_000_directory.dat). Log segment files have a state (Formatting, Preallocated, Writing, Closed, Truncated, BackedUp, RetainedFree, Free). Only log segment files with state Free can be reused. Log segment files have a fixed size although the size can vary per service. (For example, indexserver=1024MB; xsengine=8MB).

You may wish to monitor volume if, for example:

- You receive an alert about disk I/O read failure and want to see which volume has the issue and why.
- You are running out of disk space.
- You know you are having a backup issue or a replication issue and want to understand how it's affecting disk usage.

Procedure

1. Open *Disk Volume Monitor* in SAP HANA cockpit by clicking the corresponding *Monitor disk volume* link in the system *Overview*.
2. (Optional) Filter the information displayed in the chart and the table:
 - Using the arrow beside the title (top left), select a pre-defined variant, or manage and save a custom variant.
 - Select from the drop-down lists to display a specific combination of host, tenant (if applicable), volume type (data volume or log volume), service and volume ID.
3. View the information in the table:

Column	Description
Volume ID	The ID of the volume.
Service	The name of the service.
Type	Whether this is a data volume or log volume.
Size [MB]	Current size of the volume.
Used [MB]	Amount of disk space used on the host's hard disk.
Used [%]	Amount of disk space used on the host's hard disk as a percentage of the whole.

Column	Description
State	In the case of log files, whether the state is Formatting, Preallocated, Writing, Closed, Truncated, BackedUp, RetainedFree or Free. The log segment file's state indicates its availability for reuse.
Files	Number of files of the same type and state for the particular host and service.
Path	Location of the service's data and log files in the file system.
Host	The name of the host.
Tenant	The name of the tenant.

- Drill down on a row to see [Disk Volume Details](#) for a specific volume.
- Select the various tabs to move to the corresponding section of the volume details.

Section	Description
Data Volume Files	Displays the data volume file names as well as the size of each file and how much of it is currently in use, both in MB and as a percentage of its total size. Used size is the amount of data in the file. As the size of the file is automatically increased with the payload but not automatically decreased, used size and total size may be different.
Log Files	Displays log file names, total size (which, for log files, is equivalent to used size) and state. When a file is full, log entries are written to the next log segment file available.
Volume I/O Statistics	Displays aggregated I/O statistics for the volume, and, for comparison, other volumes in the system, in your choice of time periods: <ul style="list-style-type: none"> Since the service was restarted (default) Since the last manual reset You can reset the statistics collection for all volumes by selecting Reset Volume I/O Total Statistics .
Data Volume Page Statistics	Displays statistics on the data volume's pages (or blocks) broken down according to page size class. Superblocks are partitions of the data volume that contain pages of the same page size class. You can analyze how many superblocks are used for the specific size class and also how many pages/blocks are used. The fill ratio enables you to decide whether or not it makes sense to reorganize and release unnecessary superblocks, in other words, shrink the data volume.

- In the [Volume I/O Statistics](#) section, use the left-most drop-down menu to display specific statistics associated with the volume:
 - Volume Size & Time
 - Volume Configuration
 - Advanced Write Statistics
 - Advanced Read Statistics

Related Information

[Reclaim Space \[page 356\]](#)

6.6.1.5.1 Reclaim Space

You can reclaim space by reclaiming freed log segments and unused space in data volumes.

Prerequisites

Perform a backup before reclaiming space.

Context

The reclaim operation is across hosts & volumes per database & tenant. You can reclaim all log files in Free state but at least `log_preformat_segment_count` segments (by default, two) per database service will not be reclaimed. For data volumes, the reclaim is calculated by volume size – (used size * specified `percentage_of_overload_size`).

Procedure

1. Open *Disk Volume Monitor* in SAP HANA cockpit by clicking the corresponding *Monitor disk volume* link in the system *Overview*.
2. Select *Reclaim Space* from the top toolbar.
3. In the dialog, select *Reclaim (Free) log segments* and/or *Reclaim data volume*.
4. Select the *Reclaim Space* button.

6.6.1.6 Use the Statement Library to Administer Your Database

The statement library in the SAP HANA database explorer contains default system statements and can also be used to store user-defined statements.

Prerequisites

The maximum size allowed per user for all stored statements is 10 MB, uncompressed.

Only user-defined statements can be accessed by HDI containers.

You cannot use forward slashes (/) in a statement name. Otherwise, there are no character restrictions.

Context

Default system statements include all statements available in the `M_SYSTEM_INFORMATION_STATEMENTS` system view and all statements included in SAP Note 1969700.

You can edit statement contents and descriptions in the SQL console and then save them back to the statement library as a new statement, or by replacing the old version.

Create your own SQL statements for administration purposes directly in the SQL console and add them to the statement library. Alternatively, you can import statements saved as ZIP files, `.sql` files, or `.txt` files from your local computer or network file server.

The content of statements stored in the statement library is encrypted, except for the statement description; do not include sensitive information in your statement description.

The statement library is user-specific, therefore all stored statements are available for every database associated with a specific user.

Procedure

1. Create a statement and add it to the statement library.
 - a. In the catalog browser, right-click your database and click [Open SQL Console](#).
 - b. Enter your SQL statement. Including a header is optional. However, a description for your SQL statement will not show up in the statement library unless you create one in your statement header. The format for your statement description is as follows:

```
/*  
[DESCRIPTION]  
--Sample description  
*/
```

The statement description in the statement library shows the first line of the content after the `[DESCRIPTION]` section header.

- c. Click  and choose the name for the statement and the folder of the statement library where you want to store your statement.
2. Import and export SQL statements to and from the statement library.
 - a. To import a SQL statement or set of statements to the statement library, either right-click a statement library folder and click [Import](#), or click .

A file browser opens that allows you to import an individual statement in a `.sql`, `.txt`, or XML file, or a set of statements in a ZIP file.

ZIP file hierarchies are flattened in the statement library.

 - b. To export a statement from the statement library, either right-click a statement library folder and click [Export](#), or click .

Individual statements are exported as `.sql` files. Groups of statements are exported as XML files.
 3. Search for a stored statement and edit it in the SQL console.
 - a. Enter a term in the search bar in the upper right-hand corner of the [Statement Library](#) dialog () to search for in both the name and the content of all stored SQL statements.

- b. Edit a stored SQL statement by clicking its folder in the statement library and clicking .

The statement opens in the SQL console where you can edit and add it back into the statement library under the same name or a different name.

Results

You have stored user-defined statements in the statement library.

Related Information

[SAP Note 1969700](#) 

6.6.2 Monitoring in SAP HANA Studio

The SAP HANA studio has several tools for database monitoring.

- *System Monitor* 
This editor provides you with an overview of all your SAP HANA systems at a glance, including system availability and current resource usage information. From the *System Monitor* you can drill down into each individual system in the *Administration* editor.
- *Administration* 
This editor provides detailed information about resource usage, current alerts, system performance, system configuration, as well as tools for analyzing and troubleshooting issues in your system.

Related Information

[System Monitor \[page 177\]](#)

[Administration Editor \[page 359\]](#)

[Monitoring System Availability \[page 360\]](#)

[Monitoring Overall System Status and Resource Usage \[page 362\]](#)

[Monitoring Status and Resource Usage of System Components \[page 364\]](#)

[Monitoring Host Status and Auto-Failover Configuration \[page 367\]](#)

[Monitoring Alerts \[page 373\]](#)

[Monitoring Disk Space \[page 378\]](#)

[Use User-Defined SQL Statements for System Monitoring \[page 381\]](#)

[Basic Monitoring Checklist for SAP HANA Systems \[page 384\]](#)

[System and Statistics Views \[page 388\]](#)

6.6.2.1 Administration Editor

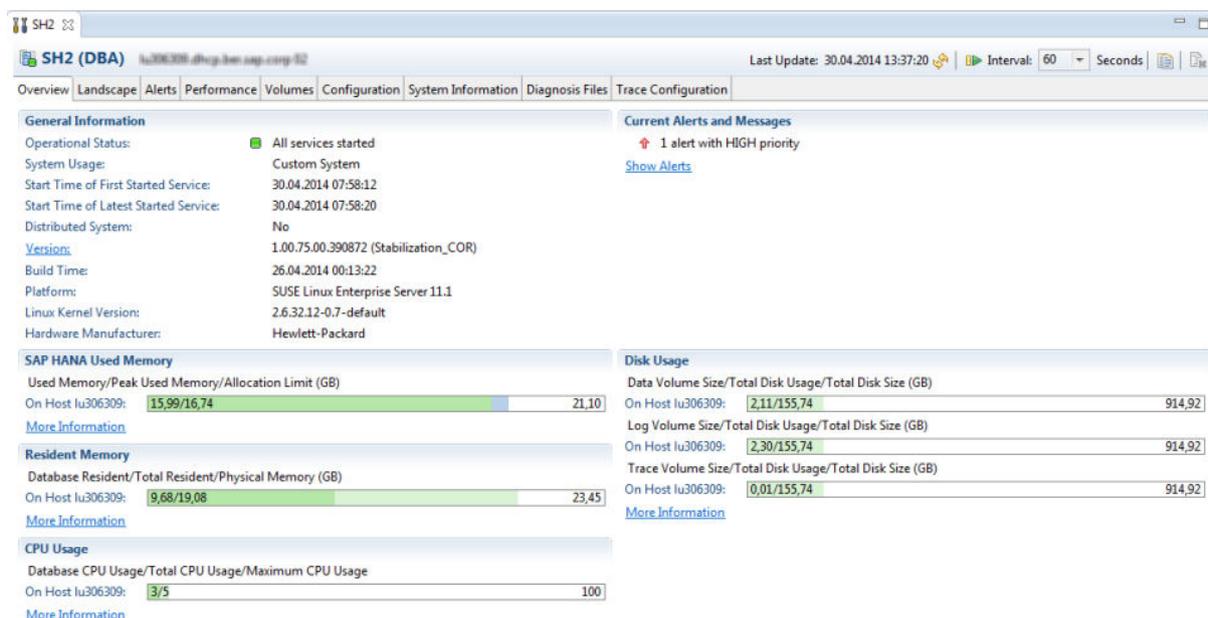
To identify problems early and avoid disruptions, you need to monitor your systems continuously. While the System Monitor provides you with an overview of all systems at a glance, the Administration editor allows you to drill down into the details of resource usage and performance for each system.

In the Administration editor, you can monitor the following:

- Overall system state and resource usage by system and host
- Status and resource usage of all system components, for example, name server, index server, and so on
- Auto-failover status and configuration of hosts in distributed systems
- Alerts issued by the system in relation to its status, performance, and resource consumption
- Disk space consumed by system processes for the various storage types (data, log, and trace)
- System performance, for example, by analyzing performance indicators such as expensive statements, running threads, and load history

Note

To open the Administration editor with read-only access to the system views and alert information, you must have either the MONITORING role or the system privilege CATALOG READ and the object privilege SELECT on the schema _SYS_STATISTICS.



The Administration Editor

Related Information

[Create and Authorize a User \[page 808\]](#)

6.6.2.2 Monitoring System Availability

The availability of an SAP HANA system is indicated by its operational status, which you can see in the *Systems* view, in the *System Monitor*, and on the *Overview* tab of the Administration editor.

The availability of an SAP HANA system is indicated by its operational status, whereby the system assumes the status of the service (*nameserver*, *indexserver*, *preprocessor*, and so on) with the most critical status.

i Note

Data and logs can only be backed up when all services that persist data are running.

The SAP HANA studio determines the status of services through an SQL connection and/or the SAP start service (*sapstartsrv*). For more information about the possible operational statuses, see *System Operational Statuses*.

You can see the operational status of a system in the SAP HANA studio in the following places:

- The *Systems* view
- The *Overview* tab of the Administration editor
- The *System Monitor*

An SAP HANA system consists of a number of services (including *indexserver*, *preprocessor*, *nameserver*, and *compileserver*).

Error Situations

Error situations can interrupt the availability of the system regardless of its operational status. A system appears in the *Systems* view with an error icon () in the following situations, for example:

- An SQL connection is not available.
- The connected user is invalid or their password has expired.
- The SAP HANA system was renamed.
- The SAP HANA license is invalid or has expired.
- The SAP HANA system is functioning as a secondary instance of your primary system, (for example, in a high availability scenario).

For more information about the nature of the error and how to resolve it, refer to the tooltip and the error log.

SAP Start Service Unreachable

An error is also indicated if *sapstartsrv* cannot be reached. If this is the case but all other services are running (their status having been determined through an SQL connection), the system itself is operational and accessible.

There are several reasons why `sapstartsrv` is not reachable. You should first check whether or not it is running. You can do this by checking the connection to the Web service in a Web browser. Enter the following URL to get the Web service description (WSDL) from the `sapstartsrv` of an SAP HANA system:

```
http://<host>:5<instance_number>13/?wsdl
```

For a secure connection, enter:

```
https://<host>:5<instance_number>14/?wsdl
```

If you receive an XML output that starts with `definitions name="SAPControl"`, then the connection is working.

In many cases, `sapstartsrv` cannot be reached because the HTTP proxy is incorrectly configured in the SAP HANA studio. To resolve this, from the main menu, choose **Window > Preferences > Network Connections** and change the value for active provider from *Native* to *Direct*.

For more information, see SAP Note 1639568.

Related Information

[Monitoring SAP HANA Systems During Stop and Start \[page 187\]](#)

[SAP Note 1639568](#)

6.6.2.2.1 System Operational Statuses

An SAP HANA system can have several operational statuses. The system assumes the status of the service with the most critical status.

Status	Description
	The status of services (and therefore the system) is unknown because a connection to the database cannot be established either through an SQL connection or <code>sapstartsrv</code> . The system is not accessible.
	All services are started. The system is operational and accessible.
	One or more services are in the process of starting or <code>sapstartsrv</code> cannot be reached.
	One or more services are not started. The system is not operational and can be accessed in diagnosis mode only.

6.6.2.3 Monitoring Overall System Status and Resource Usage

When you open the Administration editor for a particular SAP HANA system, the [Overview](#) tab provides you with a summary of the overall status of the system, as well as an overview of resource usage.

Resource usage values are presented in such a way that you can compare the SAP HANA system with the operating system as a whole. If the system is distributed across several hosts, resource usage values are aggregated across all worker hosts. An additional bar shows the host with the highest (most critical) resource usage.

The bars indicating resource usage (memory, CPU, and disk) change color (green, yellow, and red) based on configurable check thresholds.

Information Available on the Overview Tab

The following table lists the information available on the [Overview](#) tab:

Screen Area	Information Available
General Information	<ul style="list-style-type: none">• General information about the SAP HANA system, such as operational status, system usage type, whether the system has multiple hosts, the number of hosts (if distributed), and database version• The status of replication from your productive system to a secondary system This information is only available and applicable if you are operating a secondary instance of your database (for example, in a high availability scenario). If this is the case, then content from the primary or productive instance of your database is replicated to the secondary instance. More detailed information about this replication status is available on the Landscape > Secondary System Status tab.
Alerts and Messages	Priority-rated alerts and messages reported by the system

Screen Area	Information Available
Database Used Memory	<p>The following key indicators of memory usage are displayed:</p> <ul style="list-style-type: none"> • <i>Used Memory</i> The total amount of memory currently in use by SAP HANA is referred to as its used memory. • <i>Peak Used Memory</i> The <i>Used Memory</i> value is a current measurement. The <i>Peak Used Memory</i> value is the highest used memory value recorded. This is useful for keeping track of the maximum value for used memory over time. <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p>i Note</p> <p>Peak used memory is a resettable value. This can be useful if you want to establish the impact of a certain workload on memory usage. So for example, you can reset peak used memory, run the workload, and then examine the new peak used memory value. You can reset peak used memory on the ► Landscape ► Services ► tab. From the context menu, choose <i>Reset Memory Statistics</i>.</p> </div> <ul style="list-style-type: none"> • <i>Allocation Limit</i> The SAP HANA system, across its different processes, reserves a pool of memory before actual use. This pool of allocated memory is preallocated from the operating system over time, up to a predefined allocation limit, and is then efficiently used as needed by the SAP HANA database code. More memory is allocated to the pool as memory consumption increases. If the amount of memory used nears the allocation limit, SAP HANA may run out of memory if it cannot free memory.
Resident Memory	<p>Resident memory is the amount of physical memory that is actually being used from the perspective of the operating system.</p> <p>It is possible that the <i>Used Memory</i> value is lower than the <i>Database Resident</i> value if SAP HANA returns memory back to its memory pool (for example, after a temporary computation) and does not inform the operating system. This is normal.</p>
CPU Usage	The information displayed here indicates the percentage of CPU used by the SAP HANA system compared with the operating system as a whole.
Disk Usage	The information displayed here indicates disk space occupied by data, log, and trace files belonging to the SAP HANA system compared with the operating system as a whole.

Related Information

[Memory Usage in the SAP HANA Database \[page 470\]](#)

6.6.2.4 Monitoring Status and Resource Usage of System Components

To examine resource usage of an SAP HANA system in more detail, for example, to troubleshoot performance bottlenecks, you monitor its individual components or services.

On the **► Landscape ► Services ▾** tab of the Administration editor, you can see the status of the services that start when the system is started. The initial connection to the system is established by the `sapstartsrv` service. If you have a multiple-host system, the services that start depend on which components are actually installed on the instance.

For each service, detailed information about its memory consumption is available. This allows you to get a more detailed breakdown of resource usage and troubleshoot performance bottlenecks.

The available filters allow you to show and hide the information according to host and/or service. This is generally only useful if you have a multiple-host system.

Information Available on the Services Tab

The following information is displayed on *Services* sub-tab by default.

→ Tip

You can configure the view by choosing the  (*Configure Viewer*) button. For example, several additional columns are available.

Column	Description
Active	Indicates the status of the service The following statuses are possible: <ul style="list-style-type: none">• The service is started ()• The service is stopping or starting ()• The service is stopped () The daemon service displays the  icon while the host or any of its services are starting or stopping.
Host	Name of the host on which the service is running
Port	Port that the system uses for internal communication between services
Service	Service name, for example, <code>indexserver</code> , <code>nameserver</code> , <code>xsengine</code> , and so on

Column	Description
Detail	<p>Role of the host on which the service is running</p> <ul style="list-style-type: none"> • Master The host is the active master host. • <Empty> The host is an active slave host. • Standby The host is a standby host. <p>This is relevant only for distributed systems. For more detailed information, see the Hosts sub-tab.</p>
Start Time	<p>Time at which the service started</p> <p>Two of the times in this column should match the <i>Start time of first started service</i> and <i>Start time of latest started service</i> shown on <i>Overview</i> tab.</p>
Process ID	Process ID of the OS process
CPU	Bar view showing the CPU usage of the service
Memory	Bar view showing the used memory of the service in relation to physical memory available and the effective allocation limit of the service
Used Memory (MB)	Amount of memory currently used by the service
Peak Used Memory (MB)	Highest amount of memory ever used by the service
Effective Allocation Limit (MB)	Effective maximum memory pool size that is available to the process considering the current memory pool sizes of other processes
Memory Physical on Host (MB)	Total memory available on the host
SQL Port	Port through which the SQL connection to the specific service operates

Related Information

[Memory Usage in the SAP HANA Database \[page 470\]](#)

6.6.2.4.1 Reset Peak Used Memory

Resetting peak used memory allows you for example to establish the impact of a certain workload on memory usage. You can reset peak used memory in the SAP HANA studio.

Prerequisites

You have the system privilege RESOURCE ADMIN.

Context

Peak used memory is the highest recorded value for used memory. This value is useful for understanding the behavior of used memory over time and under peak loads. If you reset peak used memory and run the workload, then you can then examine the new peak used memory value.

Procedure

1. In the Administration editor, open the **▸ Landscape ▸ Services ▸** tab.
2. From the context menu, choose *Reset Memory Statistics*.
All memory statistics for all services are reset.
3. Refresh the Administration editor to see new values.

6.6.2.4.2 Memory Indicators in the SAP HANA Studio

You can see the key indicators of memory usage in various editors of the SAP HANA studio.

What	Where
Current size of used memory	<i>Overview</i> tab of the <i>Administration</i> editor
Current memory usage of individual service	▸ Landscape ▸ Services ▸ tab of the <i>Administration</i> editor
Overall peak used memory	<i>Overview</i> tab of the <i>Administration</i> editor
Peak used memory for an individual service	▸ Landscape ▸ Services ▸ tab of the <i>Administration</i> editor
Memory usage of tables broken down by schema	Execute the predefined query <i>Schema Size of Loaded Tables</i> available on the <i>System Information</i> tab of the <i>Administration</i> editor
Current effective allocation limit of a service	▸ Landscape ▸ Services ▸ tab of the <i>Administration</i> editor

Related Information

[Monitoring Overall System Status and Resource Usage \[page 362\]](#)

[Monitoring Status and Resource Usage of System Components \[page 364\]](#)

6.6.2.5 Monitoring Host Status and Auto-Failover Configuration

The SAP HANA database supports high availability in a distributed system by providing for host auto-failover. If an active host fails, for example, because of a hardware failure, standby hosts can take over and thus ensure the continued availability of the database.

You can monitor the status of individual hosts on the **► Landscape ► Hosts ▾** tab. Here, you can see all the hosts in the system, whether or not they are operational, as well as additional information about their auto-failover status and configuration.

Host roles for failover are normally configured during installation. The options available to you on the *Hosts* tab when you choose the  (*Configure Hosts for Failover Situation*) button are limited. You can only switch the configured roles of hosts; you cannot increase or decrease the number of worker hosts and standby hosts in relation to each other.

The primary reason for changing the configured roles in the *Configure Hosts for Failover Situation* dialog box is to prepare for the removal of a host. In this case, you would change the configured role of the name server host to SLAVE and the configured role of the index server host to STANDBY before stopping the database instance on the host and removing the host.

i Note

To change host configuration, you require the system privilege RESOURCE ADMIN and the object privilege EXECUTE on the procedure UPDATE_LANDSCAPE_CONFIGURATION.

Example

Typical Configuration for a Multiple-Host System

Host	Name Server (Configured Role)	Name Server (Actual Role)	Index Server (Configured Role)	Index Server (Actual Role)
Initial host	Master 1	Master	Worker	Master
1st host added	Master 2	Slave	Worker	Slave
2nd host added	Slave	Slave	Worker	Slave
3rd host added	Slave	Slave	Worker	Slave
4th host added	Slave	Slave	Worker	Slave
5th host added	Slave	Slave	Worker	Slave
6th host added	Slave	Slave	Worker	Slave
7th host added	Master 3	Slave	Standby	Standby

Information Available on the Hosts Tab

The table below lists the information available on the **► Landscape ► Hosts ▾** tab.

→ Tip

You can configure the view by choosing the  (*Configure Viewer*) button. You can also call up information about the meaning of the various statuses by choosing the  (*Display Information*) button.

Column	Description
Host	Host name
Active	<p>Indicates the status of services running on the host</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none">• YES () All services are active.• PARTIAL () Some services active.• STARTING () Some services are active, some are starting.• STOPPING () Some services are active, some are stopping• NO () No services active
Host Status	<p>Indicates whether or not the system is operational and the host's status</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none">• OK () The system is operational and the host's actual role corresponds to its configured role.• IGNORE () The system is operational. The host is configured as a standby host and is available, but not in use.• INFO () The system is operational. The host's actual role is different from its configured role.• WARNING () The system is not operational. The host will become available after start-up or failover.• ERROR () The system is not operational. The host is missing.

Column	Description
Failover Status	<p>Displays the failover status so you can see which hosts are active and which are on standby</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> • Empty Failover is neither active nor pending. • WAITING ... SEC The host has failed. The system is waiting to fail over. • WAITING The host has failed. The system is waiting for the host to restart to prevent unnecessary fail-over. • FAILOVER TO <host> The host has failed and failover to a target host is in progress. • FAILBACK TO <host> Failback to a worker host is in progress. This happens when the assigned standby host is stopped. However, there is no automatic failback while the standby host is still assigned since this would cause downtime. • FAILED Failover is not possible, for example, no further standby hosts available. For more information, see the nameserver trace.
Nameserver Role (Configured)	<p>Specifies the host's configured role as name server</p> <p>The following roles are possible:</p> <ul style="list-style-type: none"> • MASTER 1, MASTER 2, MASTER 3 When you install a distributed system, up to three hosts are automatically configured as master name servers. The configured nameserver role of these hosts is MASTER 1, MASTER 2, and MASTER 3. • SLAVE Additional hosts in your system are configured as slave name servers. The configured nameserver role of these hosts is SLAVE.
Nameserver Role (Actual)	<p>Specifies the host's actual role as name server</p> <p>The following roles are possible:</p> <ul style="list-style-type: none"> • MASTER During system start-up, one of the hosts configured as master name servers (that is, those hosts with configured name server role MASTER 1, MASTER 2, or MASTER 3) is designated to be the active master name server. The actual nameserver role of this host is MASTER. This master name server assigns one volume to each starting index server (those with actual role MASTER or SLAVE), or no volume if it is a standby host (actual indexserver role STANDBY). If this active master nameserver host fails, one of the remaining hosts configured as a master name server becomes the active master name server. • SLAVE The actual nameserver role of the remaining hosts configured as master and slave hosts is SLAVE.

Column	Description
Indexserver Role (Configured)	<p>Specifies the host's configured role as index server</p> <p>The following roles are possible:</p> <ul style="list-style-type: none"> • WORKER • STANDBY <p>When you install a distributed system, you can configure hosts either as WORKER or STANDBY index servers. A host configured as a standby index server is not used for database processing. All database processes run on the standby host, but they are idle and do not allow SQL connections.</p>
Indexserver Role (Actual)	<p>Specifies the host's actual role as index server</p> <p>The following roles are possible:</p> <ul style="list-style-type: none"> • MASTER The actual master index server is assigned on the same host as the name server with the actual role MASTER. The actual index server role of this host is MASTER. The master index server provides metadata for the other active index servers (that is, those with actual index-server role SLAVE). • SLAVE The actual index server role of remaining hosts (except those configured as standby hosts) is SLAVE. These are active index servers and are assigned to one volume. If an active index server fails, the active master name server assigns its volume to one of the standby hosts. • STANDBY The actual indexserver role of standby hosts is STANDBY. A standby host is not assigned a volume by the active master name server and it does not open an SQL port. <p>During normal operation when all hosts are available, a host with the configured role WORKER has the actual role MASTER or SLAVE, and a host with the configured role STANDBY has the actual role STANDBY. In the event of failover, the actual index server role of a host with the configured role STANDBY changes to SLAVE. The host status of the failed host changes from OK to INFO and the host status of the standby host changes from IGNORE to INFO.</p> <div style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>Failover is configured only for the name server and the index server on each host. The other components (for example, xsengine) are not configured individually as they are always failed over together with the index server.</p> </div>
Failover Group (Configured/Actual)	<p>A failover group can be defined for each host. In the event of failover, the name server tries to fail over to a host within the same group.</p>
Worker Groups (Configured/Actual)	<p>The worker groups (also referred to as a host sub-role) for the host can be set here. This is required to support heterogeneous hardware in the landscape which is required, for example, for the extension node feature.</p> <p>Worker groups may also be relevant in a single-host installation. The worker group name is a free text value, it is validated to trap illegal characters.</p>

Column	Description
Host Roles (Configured)	<p>Specifies the host's configured database or option role</p> <p>The following roles are possible:</p> <ul style="list-style-type: none"> • WORKER Worker host for database processing • STANDBY Standby host for database processing <p>Depending on your installation, the following additional host roles for SAP HANA options and capabilities may be configured:</p> <ul style="list-style-type: none"> • EXTENDED_STORAGE_WORKER Worker host for SAP HANA dynamic tiering • EXTENDED_STORAGE_STANDBY Standby host for SAP HANA dynamic tiering • ETS_WORKER Worker host for SAP HANA accelerator for SAP ASE • ETS_STANDBY Standby host for SAP HANA accelerator for SAP ASE • RDSYNC Host for SAP HANA remote data sync • STREAMING Host for SAP HANA streaming analytics
	<p>⚠ Caution</p> <p>Be aware that you need additional licenses for SAP HANA options and capabilities. For more information, see Important Disclaimer for Features in SAP HANA Platform [page 1980].</p>
	<p>SAP HANA SPS 11 includes an additional, new runtime environment for application development: SAP HANA extended application services (XS), advanced model. SAP HANA XS advanced represents an evolution of the application server architecture within SAP HANA by building upon the strengths (and expanding the scope) of SAP HANA extended application services, classic model. If the runtime platform of SAP HANA XS advanced is installed in your system, the following additional host roles are configured:</p> <ul style="list-style-type: none"> • XS_STANDBY Standby host for SAP HANA XS advanced runtime • XS_WORKER Host for SAP HANA XS advanced runtime
	<p>i Note</p> <p>Multiple host roles are not supported in production environments. However, if XS advanced runtime is installed, hosts can share multiple roles.</p>
	<p>i Note</p> <p>SAP recommends that customers and partners begin to evaluate the new capabilities of SAP HANA extended application services, advanced model with this release (SPS 11). However, please note that it is not recommended to immediately start migrating existing applications to the new capabilities.</p>

Column	Description
Host Roles (Actual)	Specifies the host's actual database or option role
Storage Partition	<p>Specifies the number of the <code>mnt000...</code> sub-directory used by the host for storing data and logs, for example, <code>1</code> if the sub-directory is <code>mnt00001</code>, <code>2</code> if it is <code>mnt00002</code>, and so on</p> <p>During installation, volumes for storing data and log files are defined. These are the directories where data and logs are stored. The default directories are:</p> <ul style="list-style-type: none"> • <code>/usr/sap/<SID>/SYS/global/hdb/data</code> for data • <code>/usr/sap/<SID>/SYS/global/hdb/log</code> for logs <p>Each active host has exactly one sub-directory beneath these directories called <code>mnt00001</code>, <code>mnt00002</code>, and so on. The next level in the file hierarchy is the actual volume, with one sub-directory for each service called <code>hdb00001</code>, <code>hdb00002</code>, and so on.</p> <p>In the event of failover, the volumes of the failed host are re-assigned to the standby host.</p>
Removal Status	<p>Indicates the status of the table redistribution operation used to move data off the index server of a host that you plan to remove</p> <p>Before you can remove an active host from a single-container system, you must move the tables on the index server of this host to the index servers on the remaining hosts in the system. You can do this by right-clicking the host and choosing <i>Remove Host...</i> Once the value in the <i>Removal Status</i> column changes to REORG FINISHED or REORG NOT REQUIRED, you can physically remove the host using the SAP HANA lifecycle management tool <code>hdblcm(gui)</code>.</p> <p>If your system is configured as a multiple-container system, you have to remove tenant-specific services first and then remove the host using the SAP HANA database lifecycle manager (HDBLCM). For more information, see <i>Remove a Service from a Tenant Database</i> in the <i>SAP HANA Tenant Databases Operations Guide</i>.</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> • <code><Empty></code> Host has not been marked for removal. • REORG PENDING A redistribution operation is required to move tables to other hosts. • REORG ACTIVE A redistribution operation is in progress. For more information, you can query the system tables <code>SYS.REORG_OVERVIEW</code> and <code>SYS.REORG_STEPS</code>. • REORG FAILED A redistribution operation was executed and failed. For more information, query the system table <code>SYS.REORG_STEPS</code>. • REORG FINISHED A redistribution operation has completed. The host can be uninstalled. • REORG NOT REQUIRED A redistribution operation not required. The host can be uninstalled.

Related Information

[High Availability for SAP HANA \[page 1080\]](#)

6.6.2.6 Monitoring Alerts

As an administrator, you actively monitor the status of the system and its services and the consumption of system resources. However, you are also alerted of critical situations, for example: a disk is becoming full, CPU usage is reaching a critical level, or a server has stopped.

The internal monitoring infrastructure of the SAP HANA database is continuously collecting and evaluating information about status, performance, and resource usage from all components of the SAP HANA database. In addition, it performs regular checks on the data in system tables and views and when configurable threshold values are exceeded, issues alerts. In this way, you are warned of potential problems. The priority of the alert indicates the severity of the problem and depends on the nature of the check and configured threshold values. For example, if 90% of available disk space is used, a low priority alert is issued; if 98% is used, a high priority alert is issued. For more information about the technical implementation of monitoring and alerting features in SAP HANA, see *The Statistics Service*.

All current unresolved alerts are summarized on the *Overview* tab of the Administration editor and displayed in detail on the *Alerts* tab. An alert remains current until the next time the relevant check is performed and the alert condition is not fulfilled, indicating that the problem situation has been resolved.

To see all alerts, on the *Alerts* tab, choose the corresponding entry from the *Show* drop-down list.

i Note

To ensure that you are seeing the latest alerts, refresh the Administration editor regularly.

Alerts are displayed according to the following time periods:

Time Period	Alerts Displayed
Last 15 minutes, last 30 minutes, last hour, last 2 hours, and today	Alerts generated in the corresponding time period are shown. If an alert was generated 10 minutes ago, it appears under all these headings.
Yesterday	Only alerts that were generated yesterday are shown.
Last week	Only alerts generated during the previous week (Sunday to Saturday) are shown.
Two weeks ago, and so on	Only alerts generated during that week are displayed.

i Note

Alerts are not rolled over into the following weeks. This enables you to compare the performance of the system over selected periods, as well as to view the alerts.

You can refine the list of displayed alerts further by specifying filters as follows:

- To filter according to a specific word in the check description, enter the word in the *Filter* field (for example, *license*).
- To filter according to additional attributes including priority and date of occurrence, choose the  *Filters* button in the toolbar on the top-right of the tab and select the required filter(s).

Detailed Alert Information

You can view the detailed information about an alert by double-clicking it. The *Alert Details* dialog box appears with information including:

- A full description of the alert
- The time stamp for this occurrence of the alert
- Information about how to resolve the alert
- A list of previous occurrences of this alert

i Note

The list is restricted to the most recent 1,000 occurrences or entries from the last 30 days.

The *Copy* button in the *Alert Details* dialog box allows you to copy the details of the alert to the clipboard, including its time(s) of occurrence. Note that only the 10 most recent occurrences are copied. Further occurrences are indicated by an ellipsis (...)

Related Information

[Alert Priorities \[page 253\]](#)

[Check Information \[page 375\]](#)

[Failing Checks \[page 393\]](#)

[Configure E-Mail Notifications for Alerts \[page 376\]](#)

[Configure Check Thresholds \[page 377\]](#)

[The Statistics Service \[page 389\]](#)

6.6.2.6.1 Alert Priorities

The priority of an alert indicates the severity of the problem and how quickly action needs to be taken.

Priority	Description
Information	Action recommended to improve system performance or stability
Low	Medium-term action required to mitigate the risk of downtime
Medium	Short-term action required (few hours, days) to mitigate the risk of downtime
High	Immediate action required to mitigate the risk of downtime, data loss, or data corruption

6.6.2.6.2 Check Information

Information about checks carried out by the system is available on the [Alerts](#) tab and in certain tables in the `_SYS_STATISTICS` schema.

Information Available on the Alerts Tab

The [Check Information](#) area at the bottom of the [Alerts](#) tab provides a full list of all the checks carried out by the system. In addition to a description of what each check does and what you need to do in the event of an alert, you can see important scheduling information.

Column	Description
ID	Check ID
Description	Check description
User Action	What you can do in the event of an alert
Max Priority	The most severe alert that was generated the last time the check was carried out For checks that consider only one object (for example, the check that determines how many days until your license expires), only one alert can be generated. This is automatically the most severe. However, for checks that consider several objects (for example, the check that determines CPU utilization in a system with multiple hosts), several alerts can be generated for one check. The most severe is recorded here. For example, if in a distributed system with 5 hosts, the CPU utilization of 2 hosts was acceptable, 2 hosts exceeded the minimum threshold value, and 1 host exceeded the medium threshold, then <code>medium</code> would be the most severe alert generated.
Last Run	When the check was last carried out
On Schedule	Whether or not the check is on schedule Even if a check is scheduled, that is there is a time and date in the Next Run column, it is important to ensure that it is actually being carried out according to schedule. Otherwise, you will not be warned about critical situations, which could have serious consequences. You can see whether or not a check is being carried out on schedule in the On Schedule column. If a check is not on schedule (<i>No</i> in the On Schedule column), then the system not performing checks as it should, and you must investigate further.
Next Run	When the check is scheduled to be carried out again
Interval	Interval at which the check is carried out in seconds

Check Information Available in the `_SYS_STATISTICS` Schema

In addition to the above information available on the [Alerts](#) tab, the following tables in the `_SYS_STATISTICS` schema provide you with further information about the default checks, their configuration and scheduled execution.

Table	Description
STATISTICS_ALERT_INFORMATION	This table describes what each check does and what to do when an alert is issued.
STATISTICS_ALERT_THRESHOLDS	This table contains the threshold values for each check. The severity level indicates the type of alert that is issued when the threshold value is reached or exceeded as follows. Severity level 1 corresponds to an information alert and level 4 a high priority alert.
STATISTICS_SCHEDULE	This table contains the scheduling information for all checks.

i Note

It also contains the scheduling information for historical data collection.

6.6.2.6.3 Configure E-Mail Notifications for Alerts

You can configure the system in such a way that you receive an e-mail when an alert condition for all or specific checks is fulfilled.

Prerequisites

- You have the system privilege CATALOG READ and the SELECT privilege on the `_SYS_STATISTICS` schema.

i Note

Both of these privileges are included in the role MONITORING.

- You have the system privilege INIFILE ADMIN.

Procedure

- In the Administration editor, choose the *Alerts* tab.
- From the tab toolbar, choose the  (*Configure Check Settings*) button. The *Configure Check Settings* dialog box appears.
- Enter the following information:
 - Sender's e-mail address
E-mail address that is entered as the email's sender
 - SMTP server
The mail server that the system sends the e-mails to

i Note

The statistics service does not support a mail server that requires additional authentication.

- Optional: SMTP Port
The default SMTP port is 25. If the configured mailserver uses a different port, enter it.
- 4. Optional: Specify the recipient(s) to whom you want an e-mail notification to be sent when an alert is generated for any check.
To do so, choose [Modify Recipients](#) and add the e-mail addresses of the users.

i Note

You can omit this step and only configure e-mail notification for specific checks (next step).

5. Specify the recipients to whom you want to an e-mail notification to be sent when an alert is generated for a specific check or checks.
 - a. Choose [Recipients Configuration for Specific Checks](#).
 - b. Select the checks for which you want to configure e-mail notification and then choose [Add Recipients](#) to add the e-mail addresses of the users to be notified.
6. Choose [OK](#) to save the configuration.

Results

The specified recipients are notified by e-mail when the system issues an alert for the relevant checks.

6.6.2.6.4 Configure Check Thresholds

For some checks performed by the system, you can configure when an alert is issued, that is the alert condition. A check can have a low, medium, and high priority threshold.

Prerequisites

- You have the system privilege CATALOG READ and the SELECT privilege on the `_SYS_STATISTICS` schema.

i Note

Both of these privileges are included in the role MONITORING.

- You have the system privilege INIFILE ADMIN.

Procedure

1. In the Administration editor, choose the *Alerts* tab.
2. From the tab toolbar, choose the  (*Configure Check Settings*) button. The *Configure Check Settings* dialog box appears.
3. Choose the *Configure Check Thresholds* tab.
4. Choose the check that you want to change and enter the threshold values.
The threshold value and unit depend on what is being measured. For example, for check 2 (disk usage), you could enter 90, 95 and 100 as the thresholds, where 90, 95 and 100 represent the percentage of disk space used.

→ Tip

Hover over a threshold value with the mouse to see information about the unit and the default value.

5. Choose *OK* when you have finished configuring the check thresholds.

Results

Alerts are generated when the system records the configured threshold values. The color of the bar views on the *Overview* tab may also change when certain thresholds are changed. For example, you change the disk space threshold from 90, 95 and 100, to 85, 90 and 95. If the disk is at 95% usage, then the bar view would change from yellow to red.

6.6.2.7 Monitoring Disk Space

To ensure that the database can always be restored to its most recent committed state, you must ensure that there is enough space on disk for data and log volumes. You can monitor disk usage, volume size, and other disk activity statistics on the *Volumes* tab of the Administration editor.

There are two views available on the *Volumes* tab for monitoring the size of volumes on disk:

- Service
- Storage type (that is data, log, and trace)

i Note

Although trace files are not stored in volumes, they are displayed on the *Volumes* tab in the *Storage* view as they consume disk space and therefore need to be monitored.

When you select a row in either view, detailed information is displayed in the lower part of the screen. In addition to size and usage information, statistics relating to the performance of read/write operations to disk are also available.

i Note

Detailed information about nameserver volumes is currently not available.

Information Available on the Volumes Tab

Service/Storage Type View of Volumes

The following table displays the information available when you select the [Service](#) view. The information shown when you select the [Storage](#) view is the same. It is simply displayed according to storage type not service. Details about the size of trace files stored on disk are also available in this view.

→ Tip

You can configure the view by right-clicking the table and selecting [Configure Table...](#) For example, several additional columns are available.

Column	Description
Service/Volume	The service host and internal port You can expand the host/port to see the storage area for data and log.
Service	The name of the service that has a data and log volume
Total Volume Size [MB]	Total size of the service's data and log volumes If you expand the host/port, you can see the size of each volume.
Data Volume Size [MB]	Current size of the service's data volume
Log Volume Size [MB]	Current size of the service's log volume
Path	Location of the service's data and log files in the file system
Storage Device ID	ID of the device on which the data and log files are stored This can be useful for checking whether or not data and log files are on the same device.
Total Disk Size [MB]	Total size of the host's hard disk
Used Disk Size [MB]	Amount of disk space used on the host's hard disk as a whole
Available Disk Size [%]	Available disk space on the host's hard disk
Volume Subpath	The subpath of a volume
Storage ID/Service	The storage ID corresponding to the service using it
Volume ID	The ID of the volume

Details View

When you select a row in either the [Storage](#) or [Service](#) view of volumes, detailed information is displayed in the lower part of the screen. In addition to size and usage information, statistics relating to the performance of read/write operations to disk are also available.

Tab Page	Description
Files	<p>This tab page displays the file name and type. It also shows the size of the file and how much of it is currently in use, both in MB and as a percentage of its total size. The relevance of used size depends on the file type as follows:</p> <ul style="list-style-type: none"> • Data files Used size is the amount of data in the file. As the size of the file is automatically increased with the payload but not automatically decreased, used size and total size may be different. • Log segment files Used size equals total size. When a file is full, log entries are written to the next log segment file available. The log segment file's state indicates its availability for reuse. For more information, see the monitoring view <code>M_LOG_SEGMENTS</code>. • Trace files Used size is zero for unused trace files and equals total size for used trace files.
Volume I/O Statistics	<p>This tab page shows aggregated I/O statistics for the volume since the service was started, for example, number of read/write requests, data throughput, total I/O time, and speed (MB/s). These figures can be useful when analyzing performance problems.</p> <p>For more information about the meaning of the individual fields, see the monitoring view <code>M_VOLUME_IO_TOTAL_STATISTICS</code>.</p>
Data Volume Superblock Statistics	<p>This tab page displays aggregated statistics on the data volume's superblocks since the service was started.</p> <p>Superblocks are partitions of the data volume that contain pages of the same page size class.</p> <p>For more information about the meaning of the individual fields, see monitoring view <code>M_DATA_VOLUME_SUPERBLOCK_STATISTICS</code>.</p>
Data Volume Page Statistics	<p>This tab page displays statistics on the data volume's pages (or blocks) broken down according to page size class. You can analyze how many superblocks are used for the specific size class and also how many pages/blocks are used. The fill ratio enables you to decide whether or not it makes sense to reorganize and release unnecessary superblocks, in other words, shrink the data volume.</p> <p>For more information about the meaning of the individual fields, see monitoring view <code>M_DATA_VOLUME_PAGE_STATISTICS</code>.</p>

Related Information

[Persistent Data Storage in the SAP HANA Database \[page 462\]](#)

6.6.2.8 Use User-Defined SQL Statements for System Monitoring

If you have your own SQL statements for monitoring purposes, you can save these on the *System Information* tab of the Administration editor for convenient repeated execution. Statements are saved in an XML file, which you can edit either directly in the studio or offline on your local file system.

Prerequisites

- The display of user-defined SQL statements on the *System Information* tab is enabled in the SAP HANA studio preferences on the **► SAP HANA > Administration** page.
- If necessary, you have changed the default name and location of the XML file to which user-defined statements are saved when you save them on the *System Information* tab.

i Note

It is possible to prepare your statements offline in an XML file and to specify this file here. The statements contained in the file then appear automatically on the *System Information* tab. However, to avoid errors, it is recommended that you create and edit statements on the *System Information* tab.

Context

For customized monitoring, it is possible to save your own SQL statements on the *System Information* tab of the Administration editor for convenient repeated execution. You can create and save individual statements directly on the *System Information* tab. Alternatively, you can import multiple statements as text or ZIP archive files from a location on your local computer or network file server. To organize large numbers of statements meaningfully, you can define a folder structure.

When you save the Administration editor, all statements, together with the defined folder structure, are saved to a single XML file and are available on the *System Information* tab of the Administration editor for all systems registered in the SAP HANA studio.

i Note

The *System Information* tab does not support prepared SQL statements. You can execute prepared statements in the SQL console.

Procedure

1. In the Administration editor, choose the *System Information* tab.
The SQL statements delivered with SAP HANA are displayed in the *System* folder.

2. Create folders for organizing your statements as required:
 - a. From the context menu, choose *New Folder*.
 - b. Enter the name and description of the folder.
3. Add user-defined statements by creating them directly or importing them from file:

Option	Description
Create a new user-defined statement	<ol style="list-style-type: none"> 1. From the context menu, choose <i>New SQL Statement</i>. 2. In the <i>User-Defined SQL Statement</i> dialog, specify a logical name and description for the statement, and then enter the statement in the space provided. 3. Save the statement.
Import user-defined statements from file	<ol style="list-style-type: none"> 1. From the context menu, choose <i>Import SQL Statements</i>. 2. Navigate to the appropriate location and select the required file(s). You can select one or more plain text files (*.txt) or ZIP archive files (*.zip) containing multiple text files. Import only flat ZIP files. Sub-directories will be ignored during import.

i Note

Statements must begin with the keyword SELECT or WITH.

The statements are added to the list of statements on the *System Information* tab. If you did an import, the individual statements contained in the text or ZIP archive files are added to the list in a new folder named *Import <timestamp>*. The name of statement is extracted from the file.

4. Save the Administration editor.
The list of statements on the *System Information* tab is saved to the XML file configured in the preferences.

Results

Statements are now available for execution on the *System Information* tab of the Administration editor for all systems registered in the SAP HANA studio.

You can edit, delete, and rearrange user-defined statements and folders.

i Note

You cannot edit or delete predefined system statements.

Related Information

[XML File Structure for User-Defined SQL Statements \[page 383\]](#)

[Execute SQL Statements in SAP HANA Studio \[page 118\]](#)

6.6.2.8.1 XML File Structure for User-Defined SQL Statements

SQL statements created in or imported into the *System Information* tab of the Administration editor are saved to a single XML file according to a defined structure.

```
<systabs version="1.0">
  <systemtables>
    <folder name="Folder 1">
      <description>Folder 1 description</description>
      <systemtable name="Statement 1">
        <description>Statement 1 description</description>
        <sql>SELECT statement</sql>
      </systemtable>
      <systemtable name="Statement 2">
        <description>Statement 2 description</description>
        <sql>SELECT statement</sql>
      </systemtable>
    </folder>
    <folder name="Folder 2">
      <description>Folder Description</description>
      <systemtable name="Statement 3">
        <description>Statement 3 description</description>
        <sql>SELECT statement</sql>
      </systemtable>
    </folder>
  </systemtables>
</systabs>
```

6.6.2.9 Basic Monitoring Checklist for SAP HANA Systems

To ensure the smooth running of your SAP HANA systems, it is important to monitor regularly operational status and key performance indicators.

Step	What to Check	What to Do
1	System availability (basic pulse check).	<p>You can verify the operational status of all your SAP HANA systems at a glance in the System Monitor and of individual systems on the Overview tab of the Administration editor.</p> <p>For full system availability, the following services must be active for each system:</p> <ul style="list-style-type: none">• <code>nameserver</code>• <code>indexserver</code>• <code>preprocessor</code> <p>You can check the status of individual services of a system on the Landscape > Services tab.</p> <p>Normally, the <code>daemon</code> service automatically restarts inactive services, but you can also do so manually by choosing Restart Missing Services from the context menu of the Landscape > Services tab.</p> <p>To investigate the reason for inactive services, consider the following actions:</p> <ul style="list-style-type: none">• On the Diagnosis Files tab, find and review the log file <code>available.log</code>, which shows whether or not the <code>daemon</code> and therefore the complete SAP HANA server was down.• On the Diagnosis Files tab, merge all diagnosis files and check the period before the service stopped (for example, the previous 30 minutes).• Check any alerts generated in the period before the service stopped.

Step	What to Check	What to Do
2	New and past alerts	<p>SAP HANA self-monitors its own status and performance and alerts you of critical situations (that is, when defined threshold values are reached or exceeded).</p> <p>New alerts appear on the Overview tab.</p> <div data-bbox="687 524 1396 674" style="background-color: #f0f0f0; padding: 10px; border-left: 2px solid #0070c0;"> <p>i Note</p> <p>For all points in this checklist, an alert is generated if a critical situation arises.</p> </div> <p>To identify potential trends and to help you troubleshoot particular issues, you should also monitor past error and high-priority alerts for specific time periods (for example, yesterday, last week, last month), as well as the frequency with which they occurred daily and weekly.</p> <p>For more information about individual alerts, refer to the details on the Alerts tab.</p>
3	Memory consumption of SAP HANA (in particular used memory and peak used memory) and memory consumption on host machines	<p>You can review the memory usage of all your systems at a glance in the System Monitor, of an individual system on the Overview tab, and of individual services on the Landscape > Services tab.</p> <p>If memory usage bars are yellow or red, consider the following actions:</p> <ul style="list-style-type: none"> • Check the error log for application or technical errors and contact SAP Support if necessary. <div data-bbox="730 1205 1396 1375" style="background-color: #f0f0f0; padding: 10px; border-left: 2px solid #0070c0;"> <p>i Note</p> <p>To open the error log, from the main menu choose Window > Show View > General > Error Log.</p> </div> <ul style="list-style-type: none"> • If there are no errors, increase available memory or reorganize your data.

Step	What to Check	What to Do
4	CPU usage	<p>You can review the CPU usage of all your systems at a glance in the System Monitor, of an individual system on the Overview tab, and of individual services on the Landscape > Services tab.</p> <p>If the CPU usage bar is yellow or red, consider the following actions:</p> <ul style="list-style-type: none"> • Check the error log for application or technical errors and contact SAP Support if necessary. <div style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p>i Note</p> <p>To open the error log, from the main menu choose Window > Show View > General > Error Log.</p> </div> <ul style="list-style-type: none"> • If there are no errors, analyze the trace files of running services. You can access these on the Diagnosis Files tab. • Increase available resources.
5	Disk usage (data volume)	<p>You can review how much disk space is being consumed by the data volume in all your systems at a glance in the System Monitor and for an individual system on the Overview tab.</p> <p>If the disk usage bar for data is yellow or red, consider the following actions:</p> <ul style="list-style-type: none"> • Reorganize your data. • Increase disk space.
6	Disk usage (log volume)	<p>You can review how much disk space is being consumed by the log volume in all your systems at a glance in the System Monitor and for an individual system on the Overview tab.</p> <p>If the disk usage bar for log files is yellow or red, consider the following actions:</p> <ul style="list-style-type: none"> • Reorganize unnecessary log files. • Verify that the last backup executed successfully.
7	Disk usage (trace files)	<p>You can review how much disk space is being consumed by trace files in all your systems at a glance in the System Monitor and for an individual system on the Overview tab.</p> <p>If the disk usage bar for trace files is yellow or red, consider the following actions:</p> <ul style="list-style-type: none"> • Switch off traces. • Delete unnecessary trace files. • Review the configured trace file rotation.

Step	What to Check	What to Do
8	Regular, successful execution of backups	<p>You can access the backup catalog, which provides information about the execution and history of data and log backups, on the <i>System Information</i> tab.</p> <p>To diagnose backup errors, refer to the files <code>backup.log</code> and <code>backint.log</code> files, which are accessible on the <i>Diagnosis Files</i> tab.</p>
9	Sufficient space for backups	Ensure that there is sufficient space at the chosen backup destination.
10	Existence of crash dump files	<p>You can check for crash dump files on the <i>Diagnosis Files</i> tab.</p> <p>If such files exist, investigate further. If necessary, contact SAP Support.</p>
11	Number of active threads and the duration of the top 5 threads	<p>You can review active threads on the ► Performance ► Threads ► tab.</p> <p>If further investigation is required, refer to the other sub-tabs of the <i>Performance</i> tab. Here you can analyze the following:</p> <ul style="list-style-type: none"> • Expensive SQL statements • Sessions • SQL performance history • Progress of long-running operations • System load history
12	Active threads with the description "call..." and the duration of such threads	"Call..." threads have a huge impact on performance. They are created on import/export of catalog objects and during data replication with using the SAP Landscape Transformation Replication Server (SAP LT).

Related Information

[Monitoring Overall System Status and Resource Usage \[page 362\]](#)

[Monitoring Alerts \[page 373\]](#)

[View Diagnosis Files in SAP HANA Studio \[page 663\]](#)

[Memory Usage in the SAP HANA Database \[page 470\]](#)

[Monitoring Status and Resource Usage of System Components \[page 364\]](#)

[Monitoring Disk Space \[page 378\]](#)

[Configure Traces in SAP HANA Studio \[page 683\]](#)

[Configure Database Trace File Rotation \[page 671\]](#)

[Monitoring and Analyzing Performance in SAP HANA Studio \[page 448\]](#)

6.6.3 System and Statistics Views

The SYS schema of the SAP HANA database contains various information about the current state of the database in its many views. Historical data is collected and stored in the views of the `_SYS_STATISTICS` schema.

System Views

The SAP HANA database provides many system views that contain important information about the database. Much of the information in these views is available in SAP HANA's administration tools. However, it can be necessary to examine the data directly as part of more detailed monitoring and performance analysis.

System views are located in the SYS schema. However, as public synonyms of all views exist, it is not necessary to specify the schema name when you query these views.

i Note

Many system views are available in two versions – one that shows the data gathered since a particular service was last started, and one that shows the data gathered since the time the view was last reset. For example, the view `M_VOLUME_IO_TOTAL_STATISTICS` shows the total read size and the total write size for each volume since a service was last started. The SQL command `ALTER SYSTEM RESET MONITORING VIEW SYS.M_VOLUME_IO_TOTAL_STATISTICS_RESET` initializes the statistics shown by this view. The view `M_VOLUME_IO_STATISTICS_RESET` now shows the statistics since the reset time.

You can access the information in a system view by querying the view directly using SQL, or opening it from the catalog using the SAP HANA database explorer or the SAP HANA studio.

→ Tip

Several predefined SQL SELECT statements on system views are available in the SAP HANA studio on the *System Information* tab of the Administration editor. These statements provide you with easy access to important system information. Double-clicking an entry in this list executes the underlying statement. To see the actual statement, from the context menu, choose *Show*.

If you have compiled your own SQL statements for monitoring purposes, you can save these statements on the *System Information* tab for convenient repeated execution.

For more information about all available system views, see the *SAP HANA SQL and System Views Reference*.

Statistics Views

The internal monitoring infrastructure of the SAP HANA database (statistics service) is continuously collecting and evaluating information about status, performance, and resource usage from all components of the SAP HANA database. This information is historicized to tables and views in the schema `_SYS_STATISTICS`. You can use these tables and views to analyze system behavior over time.

Additional System Views in the System Database

Every database has its own SYS and _SYS_STATISTICS schemas that contain information about that database only. For system-level monitoring, additional views are accessible in the system database: the M_DATABASES (SYS) view and the views in the SYS_DATABASES schema.

Related Information

[Use User-Defined SQL Statements for System Monitoring \[page 381\]](#)

[The Statistics Service \[page 389\]](#)

6.6.4 The Statistics Service

The statistics service is a central element of SAP HANA's internal monitoring infrastructure. It notifies you when critical situations arise in your systems and provides you with historical monitoring data for analysis.

- [Introduction \[page 389\]](#)
- [Technical Implementation \[page 389\]](#)
- [Data Management in the Statistics Service \[page 390\]](#)
- [The Statistics Service in Multitenant Database Containers \[page 391\]](#)

Introduction

As an SAP HANA database administrator, you need to monitor the status of the system and its services and the consumption of system resources. When critical situations arise, you need to be notified so that you can take appropriate action in a timely manner. For data center operation and resource allocation planning, you need to analyze historical monitoring data. These requirements are met by SAP HANA's internal monitoring infrastructure. A central element of this infrastructure is the statistics service.

The statistics service collects and evaluates information about status, performance, and resource consumption from all components belonging to the system. In addition, it performs regular checks and when configurable threshold values are exceeded, issues alerts. For example, if 90% of available disk space is used, a low priority alert is issued; if 98% is used, a high priority alert is issued.

Monitoring and alert information are stored in database tables in a dedicated schema (_SYS_STATISTICS). From there, the information can be accessed by administration tools, such as SAP HANA cockpit, or SAP HANA studio.

Technical Implementation

The monitoring and alerting features of the SAP HANA database are performed by the statistics service.

The **statistics service** is implemented by a set of tables and SQLScript procedures in the master index server and by the statistics scheduler thread that runs in the master name server. The SQLScript procedures either collect data (data collectors) or evaluate alert conditions (alert checkers). Procedures are invoked by the scheduler thread at regular intervals, which are specified in the configuration of the data collector or alert checker. Data collector procedures read system views and tables, process the data (for example, if the persisted values need to be calculated from the read values) and store the processed data in measurement tables for creating the measurement history.

Alert checker procedures are scheduled independently of the data collector procedures. They read current data from the original system tables and views, not from the measurement history tables. After reading the data, the alert checker procedures evaluate the configured alert conditions. If an alert condition is fulfilled, a corresponding alert is written to the alert tables. From there, it can be accessed by monitoring tools that display the alert. It is also possible to have e-mail notifications sent to administrators if an alert condition is fulfilled. Depending on the severity level of the alert, summary emails are sent with different frequency (hourly, every 6 hours, daily). You can also trigger alert checker procedures directly from monitoring tools (for example, SAP HANA studio and SAP HANA cockpit).

Data Management in the Statistics Service

The following mechanisms exist to manage the volume of data collected and generated by the statistics service:

- Configurable data retention period

The data collected by the data collectors of the statistics service is deleted after a default number of days. The majority of collectors have a default retention period of 42 days. For a list of those collectors that have a different default retention period, execute the following statement:

```
SELECT o.name, s.retention_days_default FROM
  _SYS_STATISTICS.STATISTICS_SCHEDULE s, _SYS_STATISTICS.STATISTICS_OBJECTS o
WHERE s.id = o.id AND o.type = 'Collector' and s.retention_days_default !=
42 order by 1;
```

You can change the retention period of individual data collectors with the following SQL statement:

```
UPDATE _SYS_STATISTICS.STATISTICS_SCHEDULE set
RETENTION_DAYS_CURRENT=<retention_period_in_days> where
ID=<ID_of_data_collector>;
```

→ Tip

To determine the IDs of data collectors execute the statement:

```
SELECT * from _SYS_STATISTICS.STATISTICS_OBJECTS where type = 'Collector';
```

Alert data in the `_SYS_STATISTICS.STATISTICS_ALERTS` table is also deleted by default after a period of 42 days. You can change this retention period with the statement:

```
UPDATE _SYS_STATISTICS.STATISTICS_SCHEDULE set
RETENTION_DAYS_CURRENT=<retention_period_in_days> where ID=6002;
```

- Maximum number of alerts

By default, the number of alerts in the system (that is rows in the table `_SYS_STATISTICS.STATISTICS_ALERTS_BASE`) cannot exceed 1,000,000. If this number is exceeded, the system starts deleting rows in increments of 10 percent until the number of alerts is below the maximum. To change the maximum number of alerts permitted, add a row with the key `internal.alerts.maxrows` and the new maximum value to the table `_SYS_STATISTICS"."STATISTICS_PROPERTIES`.

❖ Example

```
INSERT INTO _SYS_STATISTICS.STATISTICS_PROPERTIES VALUES  
( 'internal.alerts.maxrows', 500000 );
```

The Statistics Service in Multitenant Database Containers

In multiple-container systems, the statistics service runs as an embedded process in the (master) index server of every tenant database. Every database has its own `_SYS_STATISTICS` schema.

Monitoring tools such as the SAP HANA cockpit allow administrators in the system database to access certain alerts occurring in individual tenant databases. However, this access is restricted to alerts that identify situations with a potentially system-wide impact, for example, the physical memory on a host is running out. Alerts that expose data in the tenant database (for example, table names) are **not** visible to the system administrator in the system database.

Related Information

[Tenant Databases \[page 19\]](#)

[Using Alert and Collector Profiles \[page 391\]](#)

6.6.4.1 Using Alert and Collector Profiles

You can apply profiles to alert and collector values.

The internal monitoring infrastructure of the SAP HANA database is continuously collecting and evaluating information about status, performance, and resource usage from all components of the SAP HANA database. You can control the alert and collector behavior by applying profiles.

Available Profiles

Each profile affects the system performance in a different way. Select the profile that suits your system best. As each profile may change the activity status, the frequency of execution and the retention time (of collector data) of alerts and collectors of the statistics service, the amount of resources consumed by the statistics service may change accordingly. All profiles may change the value of the parameter

`retention_days_default` which controls for how many days collector data is stored. Additionally, the profile `HXE` reduces the number of enabled alerts and collectors.

The following profiles are available:

Profile	Description
S	This profile applies default values multiplied by 0.5. Resulting values are rounded up, if necessary.
M	No changes are applied to the default alert and collector values. This is the default profile.
L	This profile applies default values multiplied by 1.5. Resulting values are rounded up, if necessary.
HXE	This profile disables several alerts and collectors, and sets the default values for all collectors to 1. This is the default profile for SAP HANA Express systems.

Switching Between Profiles

The following mechanisms exist to manage alert and collector profiles:

- View the currently enabled profile with the following statement:

```
SELECT VALUE FROM _SYS_STATISTICS.STATISTICS_PROPERTIES WHERE KEY =  
'internal.sizing.profile'
```

- Enable a profile with the following statement:

```
UPDATE _SYS_STATISTICS.STATISTICS_PROPERTIES SET VALUE = '<S, M, L, HXE>'  
WHERE KEY = 'internal.sizing.profile'
```

Customized values for `DEFAULT_VALUE` in the table `_SYS_STATISTICS.STATISTICS_ALERT_THRESHOLDS` and `STATUS`, `INTERVALLENGTH`, `RETENTION_DAYS_DEFAULT` in the table `_SYS_STATISTICS.STATISTICS_SCHEDULE` are stored with each profile.

SAP HANA Express Profile (HXE)

The SAP HANA Express profile disables a number of alerts and sets the default values for all collectors to 1.

The following alerts remain active after enabling the HXE profile:

- Alert_Check_Inactive_Services
- Alert_Check_Restarted_Services
- Alert_Check_Service_Allocation_Limit
- Alert_CrashDump_Files
- Alert_Internal_Disk_Full_Events
- Alert_Internal_Events
- Alert_License_Expiring

- Alert_Lock_Wait_Time_Out
- Alert_Long_Running_Statements
- Alert_Mon_TraceFileSize
- Alert_Password_Expiration
- Alert_RTE_Dump_Files
- Alert_Summary_Email_All
- Alert_Summary_Email_High
- Alert_Summary_Email_Medium_High

The following collectors remain active after enabling the HXE profile:

- Collector_Global_Internal_Events
- Collector_Host_Column_Tables_Part_Size
- Collector_Host_Heap_Allocators
- Collector_Host_Service_Memory
- Collector_Tel_Disk_Usage
- Collector_Tel_Feature_Usage
- Collector_Tel_Host_Information
- Collector_Tel_Inifile_Contents
- Collector_Tel_Licenses
- Collector_Tel_Out_Of_Memory_Events
- Collector_Tel_System_Overview

Related Information

[Monitoring Alerts \[page 373\]](#)

[Monitoring and Analyzing Performance in SAP HANA Studio \[page 448\]](#)

6.6.4.2 Failing Checks

The alerting mechanism of the SAP HANA database relies on the regular execution of checks. If a check fails to execute, it is important to investigate the reason why. Otherwise, you may not be warned about potentially critical situations. Checks often fail due to a shortage of system resources.

If a check fails to execute, an alert is issued indicating that there is an internal statistics service problem. You can also see whether individual checks have stopped running on schedule in the *Check Information* area of the *Alerts* tab. As long as a check is not being executed, it cannot alert you about potentially critical situations.

Alerting Mechanism	Response
Statistics service	<p>A check is disabled the first time it fails to execute. It remains disabled for a specific length of time before it is automatically re-enabled. This length of time is calculated based on the values in the following columns of the table STATISTICS_SCHEDULE (_SYS_STATISTICS):</p> <ul style="list-style-type: none"> INTERVALLENGTH SKIP_INTERVAL_ON_DISABLE <p>Once INTERVALLENGTH x SKIP_INTERVAL_ON_DISABLE has elapsed, the check is re-enabled. The default values for all checks are such that failed checks remain disabled for 1 hour.</p> <p>The system determines the status of every check and/or whether the time to re-enablement has elapsed every 60 seconds.</p> <p>You can control the time to re-enablement by changing the value in the column SKIP_INTERVAL_ON_DISABLE.</p> <p>You can also re-enable the check manually.</p>

i Note

The behavior described above also applies to the data collectors of the statistics service.

Related Information

[Switch Alerting Off/On \[page 350\]](#)

[SAP Note 1991615](#)

6.7 Managing and Monitoring the Performance of SAP HANA

In addition to the tools in the SAP HANA cockpit and the SAP HANA studio, other resources are introduced here that are available to you to help improve the performance of your system.

Monitoring

Monitoring past and current information about the performance of the SAP HANA database is important to prevent performance issues and for root-cause analysis of problems. The SAP HANA cockpit and the SAP HANA studio provide a number of monitoring tools; the Performance Monitor of the SAP HANA cockpit is particularly useful for analysis as it shows side-by-side visual displays of both system performance and the workload currently being applied.

Caching

Caching is used widely in SAP HANA as a strategy to improve performance by re-using queried data rather than re-reading and processing the data every time it is requested. Administration and analysis tools for

working with the Plan Cache for query optimization are available in the SAP HANA Cockpit and Studio. In addition to this, the following configurable applications of caching query results in SAP HANA are available:

The query result cache applies to column store tables and views and offers most potential to improve performance in situations where data is predominantly read; this is because updates to the base tables invalidate this type of cache. Configuration options for this cache include defining a maximum memory budget for the cache and defining a white list of tables or views to which caching is applied. This is described in detail in SAP Note 2014148 *Guidelines for Using the Query Result Cache*.

The static result cache (sometimes referred to as *cached views*) and the dynamic result cache are further applications of caching. The static result cache is created for a specific view and remains valid for the duration of a user-defined retention period. The dynamic result cache is similar but does not have a retention period; it guarantees transactional consistency by maintaining delta records of all changes applied to the underlying table. Details of the static and dynamic result cache can be found in the *SAP HANA Troubleshooting and Performance Analysis Guide*.

Hints

Hints are instructions for the SAP HANA database server which influence the way a database request is processed. They are typically used to optimize SAP HANA performance or memory consumption and have no effect on the result set of the request. Predefined hints for various purposes are delivered with SAP HANA, you can list these from the HINTS system view; refer to the *SAP HANA SQL and System Views Reference Guide* for more details. You can also create user-defined hints and apply these to a select statement. If a query has both a user-defined and a system-defined hint, the user hint is used in preference.

For convenience you can associate queries with specific hints so that the hint is always applied to the query at runtime. This can be done firstly in the Statement Hints app of SAP HANA Cockpit (see also STATEMENT_HINTS system view), and as an alternative to linking a hint to a select statement you can pin a hint to a specific execution plan in the SQL plan cache. Pinning is applied using the ALTER SYSTEM PIN SQL PLAN CACHE ENTRY instruction and uses the execution plan PLAN_ID value to link to a hint. Refer to the *SAP HANA SQL and System Views Reference Guide* and to *SAP Note 2400006 FAQ: SAP HANA Statement Hints* for details.

Many hints are available for use with cached data or data replicas to provide 'hint-based routing' functionality so that administrators have control over exactly which source is used when a choice of data sources is available. Here, in this Performance chapter of the *SAP HANA Administration Guide* we also include a section on the configurable hint classes which are available, see *Using Hints to Query Data Snapshots* for details.

Capture and Replay

SAP HANA capture and replay allows you to capture the workload of a production system and to replay the captured workload on a target system. This can help you evaluate potential impacts on performance or stability after a change in hardware or software configuration.

Related Information

[Performance: Using Hints to Query Data Snapshots \[page 457\]](#)

[Managing Statement Hints \[page 448\]](#)

[Capturing and Replaying Workloads \[page 410\]](#)

[SAP Note 2400006](#) 

6.7.1 Monitoring, Managing, and Analyzing Performance in SAP HANA Cockpit

You can manage, monitor, analyze, and improve the performance of the database using the SAP HANA cockpit.

Related Information

[Monitoring Performance in SAP HANA Cockpit \[page 396\]](#)

[Managing Performance in SAP HANA Cockpit \[page 409\]](#)

[Analyzing Performance in SAP HANA Cockpit \[page 435\]](#)

[Improving Performance in SAP HANA Cockpit \[page 446\]](#)

6.7.1.1 Monitoring Performance in SAP HANA Cockpit

Monitoring past and current information about the performance of the SAP HANA database is important for root-cause analysis and the prevention of future performance issues.

You can use the following tools to monitor fine-grained aspects of system performance in the SAP HANA cockpit:

- Use the *Performance Monitor* to visually analyze historical performance data across a range of key performance indicators related to memory, disk, and CPU usage.
- Use *Threads* to monitor the longest-running threads active in your system. You can use it to see, for example, how long a thread is running, or if a thread is blocked for a prolonged period.
- Use the *Sessions* tile to monitor all sessions in your landscape.
- Use the *Statements Monitor* to analyze the current most critical statements running in the database.
- Use *Expensive Statements* to analyze individual SQL queries whose execution time was above a configured threshold.
- Use the *SQL plan cache* to get an insight into the workload of the SAP HANA database as it lists all statements currently cached in the SAP HANA database.
- Use the *Blocked Transactions* to monitor the details of transactionally blocked threads.

Related Information

[Monitoring and Analyzing with the Performance Monitor \[page 397\]](#)

[Monitoring and Analyzing Threads \[page 402\]](#)

[Monitoring and Analyzing Sessions \[page 406\]](#)

[Monitoring and Analyzing with the Statements Monitor \[page 406\]](#)

[Monitoring and Analyzing Expensive Statements \[page 407\]](#)

[Monitoring and Analyzing Statements with SQL Plan Cache \[page 408\]](#)

[Monitoring Blocked Transactions \[page 408\]](#)

6.7.1.1.1 Monitoring and Analyzing with the Performance Monitor

Analyzing the performance of the SAP HANA database over time can help you pinpoint bottlenecks, identify patterns, and forecast requirements. Use the *Performance Monitor* to visually analyze historical performance data across a range of key performance indicators related to memory, disk, and CPU usage.

Open the *Performance Monitor* by clicking the chart or the *Show all* link on the *Memory Usage*, *CPU Usage*, or *Disk Usage* tile on the homepage of the SAP HANA cockpit.

The *Performance Monitor* opens displaying the load graph for the selected resource: CPU, disk, or memory. The load graph initially visualizes resource usage of all hosts and services listed on the left according to the default KPI group of the selected resource.

You can customize the information displayed on the load graph, for example:

- Define the monitored time frame by entering your desired dates or selecting from *Presets*
- Set the automatic refresh rate
- Use the *Add Chart* button to create custom charts displaying the host and services selection, and selected KPIs
For a list of all available KPIs, see *Key Performance Indicators*.
- Zoom into a specific time on a graph by brushing across the desired selection on the load graph directly. Click on the zoom in button on upper right corner of the highlighted area. Select *Undo* in the header toolbar to zoom out again
- Compare the performance of your selected KPIs at different times using the *Performance Comparison* page. For more information, see *Compare Performance*.
- In the *Settings* menu, customize your graphs by including hosts and services as well as additional KPIs in the *Charts* tab. In the *Alerts* tab, configure alerts according to category and priority status.

Related Information

[Monitoring and Analyzing with the Performance Monitor \[page 397\]](#)

[Key Performance Indicators \[page 249\]](#)

[Compare Performance \[page 400\]](#)

[Collecting Performance Monitor Data for SAP Support \[page 401\]](#)

[Import Performance Monitor Data \[page 401\]](#)

[Export Performance Monitor Data \[page 402\]](#)

6.7.1.1.1 Key Performance Indicators

The *Performance Monitor* allows you select a range of host-level and service-level KPIs to analyze historical performance data of the SAP HANA database.

Host KPIs

KPI	Description
CPU	CPU used by all processes related to the operating system (OS)
Database resident memory	Physical memory used by all SAP HANA database processes
Total resident memory	Physical memory used by all OS processes
Physical memory size	Total physical memory
Database used memory	Memory used by all SAP HANA database processes
Database allocation limit	Memory allocation limit for all SAP HANA database processes
Disk used	Disk space used by data, log, and trace files belonging to the SAP HANA database
Disk size	Total disk size
Network in	Bytes read from the network by all processes
Network out	Bytes written to the network by all processes
Swap in	Bytes read from swap memory by all processes
Swap out	Bytes written to swap memory by all processes

Services KPIs

KPI	Description
CPU	CPU used by the database process
System CPU	CPU used by the database process relative to the operating system
Memory used	Memory used by the database process
Memory allocation limit	Effective allocation limit of the database process
Handles	Number of open handles in the index server process
Ping time	Indexserver ping time including nsWatchdog request and collection of service-specific KPIs
Swap in	Bytes read from swap by the process
Open connections	Number of open SQL connections
Open transactions	Number of open SQL transactions

KPI	Description
Blocked transactions	Number of blocked SQL transactions
Statements	Number of finished SQL statements
Active commit ID range	Range between newest and oldest active commit ID
Pending session request count	Number of pending requests
Active versions	Number of active MVCC versions
Acquired record locks	Number of acquire record locks
Read requests	Number of read requests (selects)
Write requests	Number of write requests (insert, update, and delete)
Merge requests	Number of merge requests
Column unloads	Number of table and column unloads
Active threads	Number of active threads
Waiting threads	Number of waiting threads
Total threads	Total number of threads
Active SqlExecutors	Total number of active SqlExecutor threads
Waiting SqlExecutors	Total number of waiting SqlExecutor threads
Total SqlExecutors	Total number of SqlExecutor threads
Data write size	Bytes written to data area
Data write time	Time used for writing to data area
Log write size	Bytes written to log area
Log write time	Time used for writing to log area
Data read size	Bytes read from data area
Data read time	Time used for reading from data area
Log read size	Bytes read from log area
Log read time	Time used for reading from log area
Data backup write size	Bytes written to data backup
Data backup write time	Time used for writing to data backup
Log backup write size	Bytes written to log backup
Log backup write time	Time used for writing to log backup
Mutex Collisions	Number of collisions on mutexes
Read/Write Lock Collisions	Number of collisions on read/write locks

Related Information

[Memory Usage in the SAP HANA Database \[page 470\]](#)

6.7.1.1.2 Compare Performance

Use [Performance Comparison](#) to examine the performance of your selected KPIs at different time intervals.

Procedure

1. In the [Performance Monitor](#), select the KPIs you want to compare.
2. Brush across the desired time selection on the load graph directly.

This selection will make up the main chart that you can contrast to any additional charts you create on the [Performance Comparison](#) page.

3. Click on the comparison button on the upper right corner of the highlighted area.
4. The [Performance Comparison](#) page opens, displaying the KPIs as well as hosts and services that were selected in the [Performance Monitor](#).
5. Optional: You can add or remove KPIs by clicking the [Refine KPIs](#) button in the header toolbar and making your selection. You can also adjust hosts and services.

Optional: You can adjust the time range of the chart by selecting the desired start and end of the monitored time interval, or choosing from [Presets](#) in the header toolbar.

6. Add an additional chart for comparing performance at different time intervals by selecting the [Add a chart to compare](#) link on the bottom of the screen or by clicking the [Add chart](#) button in the header toolbar.

A selection of preset time intervals to choose from opens. Once you have made your choice, the additional chart displaying that time range appears.

i Note

The [Add a chart to compare](#) link is only available for the first additional chart, any other chart must be added by using the [Add Chart](#) button in the header toolbar.

7. Optional: Per default, the monitored time interval is defined via a range. To choose a time interval that is dynamically adjusted to the time interval of the main chart, click on the [Relational](#) button above the respective chart and make your time interval selection.
8. Optional: Update the chart by pressing the [Update](#) button above the respective chart.
9. Optional: You can bookmark a time range in a load chart to easily refer to it in the future.

Highlight a time range on the desired chart, click the navigation icon on the top right corner of the highlighted area, and choose [Bookmark Selection](#).

The highlighted area changes color to indicate that a bookmark has been set. Above the chart containing bookmarks, there is a link with the number of bookmarks contained in the chart. It lists the bookmarked time range as well as the bookmark selection date. Clicking it highlights and navigates to the bookmarked time range on the chart.

It is possible to name the bookmark by clicking the navigation icon on the highlighted area and selecting [Add Description](#). The description is displayed in the bookmark list above the chart.

To modify the description, click on the navigation icon. You can also delete the bookmark through the navigation icon or by clicking the trash bin.

10. Navigate to the [Performance Monitor](#) or [Workload Analysis](#) page by highlighting a time range, clicking on the navigation icon and making your selection.

6.7.1.1.1.3 Collecting Performance Monitor Data for SAP Support

To help SAP Support analyze and diagnose problems with your system, you can collect a snapshot of the performance monitor data from your system into a zip file. You can trigger the collection of diagnosis information from the SAP HANA cockpit.

Related Information

[Import Performance Monitor Data \[page 401\]](#)

[Export Performance Monitor Data \[page 402\]](#)

6.7.1.1.1.3.1 Import Performance Monitor Data

To analyze and diagnose problems with the SAP HANA database, you can import performance monitor data from a zip file into the SAP HANA cockpit.

Procedure

1. Open the [Performance Monitor](#) by clicking the chart or the [Show all](#) link on the [Memory Usage](#), [CPU Usage](#), or [Disk Usage](#) tile on the homepage of the SAP HANA cockpit.
The [Performance Monitor](#) opens displaying the load graph for the selected resource: CPU, disk, or memory.
2. Select [Import](#) and select the file containing the performance monitor data set that you want to import.
Click [Import](#) to confirm your selection.

The system imports the performance monitor data set from the zip file. This may take some time and runs in the background.

Once the performance monitor data is available, it is displayed in the list of [Performance Monitor Data Sets](#).

Next Steps

You can open a performance monitor data set by clicking the corresponding entry under [Performance Monitor Data Sets](#). The [Performance Monitor](#) opens and displays the KPI data stored inside the data set. Use the

[Performance Monitor](#) to visually analyze historical performance data across a range of key performance indicators related to memory, disk, and CPU usage.

Related Information

[Monitoring and Analyzing with the Performance Monitor \[page 397\]](#)

6.7.1.1.3.2 Export Performance Monitor Data

To help SAP Support analyze and diagnose problems with the SAP HANA database, you can export performance monitor data into a zip file, which you can then download and, for example, attach to a support message.

Procedure

1. Open the [Performance Monitor](#) by clicking [Show all](#) on the [Memory Usage](#), [CPU Usage](#), or [Disk Usage](#) tile on the homepage of the SAP HANA cockpit.
The [Performance Monitor](#) opens to a display of the load graph for the selected resource: CPU, disk, or memory.
2. Select [Export All](#) in the footer bar to export the CPU, disk, or memory KPI data.

Click [Export](#) in the [Export All](#) dialog. The system collects the relevant information and saves it to a zip file. This may take some time and runs in the background.

Once the collection is available, you can download it by clicking the download button. It will be saved to the download directory of your browser on your client.

Related Information

[Monitoring and Analyzing with the Performance Monitor \[page 397\]](#)

6.7.1.1.2 Monitoring and Analyzing Threads

Use [Threads](#) to monitor the longest-running threads active in your system. You can use it to see, for example, how long a thread is running, or if a thread is blocked for a prolonged period.

Analyzing the threads running in the SAP HANA database can be helpful when analyzing the current system load.

You can identify which statements or procedures are being executed and at what stage they are, who else is connected to the system, and if there are any internal processes running as well.

The *Threads* tile provides information about the number of currently active and blocked threads in the database.

Open the *Threads* tile by clicking either the number of active threads or blocked threads on the tile.

The *Threads* page allows you to monitor the longest-running threads in your current system. You can retrieve more information or customize what is being displayed, for example:

- Filter threads by host, service, and thread type
- Choose the sorting order by checking the *Group and Sort* box and selecting the sorting parameters
- See the call stack information on your chosen thread
- Define columns and choose the parameters you want information on

When you have selected a thread, you can *Navigate To* the *Sessions* or *Blocked Transactions* page for the thread with the same connection ID.

If a thread is in a blocked transaction or is using an excessive amount of memory, you can cancel the operation executing the thread by clicking *Cancel Operations* in the footer toolbar.

Related Information

[Thread Details \[page 403\]](#)

6.7.1.1.2.1 Thread Details

The *Threads* tile provides you with detailed information about the 1000 longest-running threads currently active in the database.

i Note

Not all of the columns listed below are visible by default. You can add and remove columns in the table personalization dialog, which you open by clicking the personalization icon in the table toolbar.

Thread Information

The table below lists the information available for threads.

Detail	Description
Blocking Transaction	Blocking transaction
Duration (ms)	Duration (ms)

Detail	Description
Host	Host name
Port	Internal port
Service	Service name
Hierarchy	Thread grouping information. Filled with Connection ID/ Update Transaction ID/Transaction ID or left empty for inactive threads
Connection ID	Connection ID
Thread ID	Thread ID
Calling	The thread or service which the thread calls
Caller	The thread or service which called this thread
Thread Type	Thread type
Thread Method	Thread method
Thread Detail	Thread detail
User	User
Application User	Application user name
CPU Time	CPU time of thread
Cumulative CPU Time	CPU time of thread and associated children
Transaction ID	Transaction ID
Update Transaction ID	Update transaction ID
Thread Status	Thread state
Connection Transaction ID	Transaction object ID
Connection Start Time	Connected Time
Connection Idle Time (ms)	Time that the connection is unused and idle
Connection Status	Connection Status: 'RUNNING' or 'IDLE'
Client Host	Host name of client machine
Client IP	IP of client machine
Client PID	Client Process ID
Connection Type	Connection type: Remote, Local, History (remote), History (local)
Own Connection	Own connection: TRUE if own connection, FALSE if not
Memory Size per Connection	Allocated memory size per connection
Auto Commit	Commit mode of the current transaction: TRUE if the current connection is in auto-commit mode, FALSE otherwise

Detail	Description
Last Action	The last action done by the current connection: ExecuteGroup, CommitTrans, AbortTrans, PrepareStatement, CloseStatement, ExecutePrepared, ExecuteStatement, FetchCursor, CloseCursor, LobGetPiece, LogPutPiece, LobFind, Authenticate, Connect, Disconnect, ExecQidltab, CursorFetchltab, InsertIncompleteltab, AbapStream, TxStartXA, TxJoinXA
Current Statement ID	Current statement ID
Current Operator Name	Current operator name
Fetch Record Count	Sum of the record count fetched by select statements
Sent Message Size (Bytes)	Total size of messages sent by the current connection
Sent Message Count	Total message count sent by the current connection
Received Message Size (Byte)	Total size of messages/transactions received by the current connection
Received Message Count	Total message/transaction count received by the current connection
Creator Thread ID	Thread ID who created the current connection
Created By	Engine component that created the connections: Session, Planning, Repository, CalcEngine, Authentication, Table Exporter, Loader, LLVM, JSVM, IMS Search API, OLAP Engine, Mergedog, Ping Status, Name Server, Queue Server, SQL Stored Procedure, Authorization, TrexViaDbssl from ABAP, HybridTable Reorganizer, Session external
Is Encrypted	Encrypted: TRUE if the secure communication is enabled (SSL enabled), FALSE, otherwise
Connection End Time	The time when the connection is closed for history connections
Blocked Update Transaction ID	Write transaction ID of the write transaction waiting for the lock
Blocking Transaction ID	Transaction object ID of the transaction holding the lock
Thread ID of Lock Owner	Connection ID associated with the blocked write transaction
Blocking Update Transaction ID	Write transaction ID of the write transaction holding the lock
Transactional Lock Type	Transactional lock type
Transactional Lock Mode	Transactional lock mode
Lock Wait Component	Waiting for lock component
Lock Wait Name	Waiting for lock ID
Timestamp of Blocked Transaction	Timestamp of the blocked transaction
Waiting Record ID	ID of the record on which the lock is currently placed
Waiting Object Name	Name of the object on which the lock is currently placed
Waiting Object Type	Type of the object on which the lock is currently placed

Detail	Description
Waiting Schema Name	Name of the schema on which the lock is currently placed

6.7.1.1.3 Monitoring and Analyzing Sessions

Use the [Sessions](#) tile to monitor all sessions in your landscape.

Analyzing the sessions connected to your SAP HANA database helps you identify which applications or which users are currently connected to your system, as well as what they are doing in terms of SQL execution.

The [Sessions](#) tile displays the number of active and total sessions.

Open the [Sessions](#) tile.

The [Sessions](#) page allows you to monitor all sessions in the current landscape. You can see the following information:

- Active/inactive sessions and their relation to applications
- Whether a session is blocked and, if so, which session is blocking it
- The number of transactions that are blocked by a blocking session
- Statistics such as average query runtime and the number of DML and DDL statements in a session
- The operator currently being processed by an active session

To support monitoring and analysis, you can perform the following actions on the [Sessions](#) page:

- Cancel a session by choosing [Cancel Sessions](#)
- Save the data sets as a text or html file by choosing the [Save As...](#) button.

6.7.1.1.4 Monitoring and Analyzing with the Statements Monitor

Use [Monitor Statements](#) to analyze the current most critical statements running in the database.

Analyzing the current most critical statements running in the SAP HANA database can help you identify the root cause of poor performance, CPU bottlenecks, or out-of-memory situations. Enabling memory tracking allows you to monitor the amount of memory used by single statement executions.

The [Monitor Statements](#) tile displays the number of long-running statements and long-running blocking situations currently active in the database. Statements are ranked based on a combination of the following criteria:

- Runtime of the current statement execution
- Lock wait time of the current statement execution
- Cursor duration of the current statement execution

Open the [Monitor Statements](#) page by clicking either the long-running statements or long-running blocked statements on the tile.

The [Monitor Statements](#) page allows you to analyze the most current statements running in the database. You can see:

- The 100 most critical statements, listed in order of the longest runtime
- The full statement string and ID of the session in which the statement is running
- Application, application user, and the database running the statement
- Whether a statement is related to a blocking transaction

To support monitoring, you can perform these actions on the [Monitor Statements](#) page:

- If a statement is in a blocked transaction or using an excessive amount of memory, cancel the session the statement is running in (or the blocked session) by clicking [Cancel Session](#) in the footer toolbar
- Access information about the memory consumption of statements by clicking [Enable Memory Tracking](#) in the footer toolbar
- Set up or modify workload classes by clicking a statement's [Workload Class](#) name. Choose [New](#) to create a new workload class or [Existing](#) to select a workload class from a list, then fill in the fields.

Related Information

[Setting a Memory Limit for SQL Statements \[page 633\]](#)

[Create a Workload Class \[page 650\]](#)

[Create a Workload Class Mapping \[page 651\]](#)

6.7.1.1.5 Monitoring and Analyzing Expensive Statements

Use [Expensive Statements](#) to analyze individual SQL queries whose execution time is above a configured threshold.

Analyzing expensive statements can help you understand why they exceed duration thresholds.

You can find the [Monitor expensive Statements](#) link in the [Monitoring](#) link list on the SAP HANA cockpit.

The expensive statements trace records information about the expensive statements for further analysis and displays it on the [Expensive Statements Trace](#) page.

To support monitoring and analysis, you can perform these actions on the [Expensive Statements Trace](#) page:

- The expensive statements trace is deactivated by default. Activate and configure it by selecting [Configure Trace](#) in the footer bar.
- Define the monitored date.
- Filter expensive statements, refresh the list, choose the sorting parameter, and filter by parameter.
- Save the data sets as a text or html file by choosing the [Save As...](#) button.
- Configure the threshold parameters by clicking the [Configure Trace](#) button and entering information on the [Configure Expensive Trace](#) page.
- Open an expensive statement with the SQL analyzer by clicking [More](#) next to the statement string.
- Set up or modify workload classes by clicking a statement's [Workload Class Name](#). Choose [New](#) to create a new workload class or [Existing](#) to select a workload class from a list, then fill in the fields.

Related Information

[Expensive Statements Trace \[page 676\]](#)

[Create a Workload Class \[page 650\]](#)

[Create a Workload Class Mapping \[page 651\]](#)

6.7.1.1.6 Monitoring and Analyzing Statements with SQL Plan Cache

Use the *SQL plan cache* to get an insight into the workload of the SAP HANA database as it lists all statements currently cached in the SAP HANA database.

Analyzing all statements currently cached in the SAP HANA database can help you identify statement hashes, as well as if a statement has been correctly cached.

You can find the *SQL plan cache* link in the *Administration* link list on the SAP HANA cockpit.

Technically, the plan cache stores compiled execution plans of SQL statements for reuse, which gives a performance advantage over recompilation at each invocation. For monitoring reasons, the plan cache keeps statistics about each plan, for instance number of executions, min/max/total/average runtime, and lock/wait statistics. Analyzing the plan cache is very helpful as one of the first steps in performance analysis because it gives an overview about what statements are executed in the system.

i Note

Due to the nature of a cache, seldom-used entries are evicted from the plan cache.

The SQL plan cache is useful for observing overall SQL performance as it provides statistics on compiled queries. You can get insight into frequently executed queries and slow queries with a view to finding potential candidates for optimization.

To support monitoring and analysis, you can perform the following actions on the *SQL Plan Cache* page:

- To open an SQL statement with the SQL analyzer, you can do so by clicking *More* next to the statement string.
- Save the data sets as a text or html file by choosing the *Save As...* button.
- The collection of SQL plan cache statistics is enabled by default, but you can disable it by choosing *Configure*.

6.7.1.1.7 Monitoring Blocked Transactions

Use *Blocked Transactions* to monitor transactionally blocked threads. You can use it to see, for example, what transaction is blocking a thread, the type of lock held, and the owner of the lock.

Blocked transactions are transactions that are unable to be processed further because they need to acquire transactional locks (record or table locks) that are currently held by another transaction. Transactions can also be blocked while waiting for other resources such as network or a disk (database or metadata locks).

Analyzing the blocked transactions in SAP HANA database can be helpful when analyzing the current system load, as transactionally blocked threads can impact application responsiveness.

You can find the [Open Blocked Transactions](#) link in the [Monitoring](#) list on the SAP HANA cockpit.

The [Blocked Transactions](#) feature provides information on the number of currently blocked threads in the database.

To support monitoring and analysis, you can perform the following actions on the [Blocked Transactions](#) page:

- Filter transactions with the help user-defined keywords
- Select to hide own or idle sessions
- Customize the blocked transaction columns to show only desired parameters
- Click on a blocked transaction and select [Navigate To...](#) on the bottom right of the screen to jump to [Threads](#) or [Sessions](#) with the same connection ID.

6.7.1.2 Managing Performance in SAP HANA Cockpit

You can manage the performance of SAP HANA using SAP HANA capture and replay.

Use SAP HANA capture and replay to detect, analyze, or verify any potential issues before applying changes or upgrades, such as:

- Hardware change
- SAP HANA revision upgrade
- SAP HANA ini file change
- Table partitioning change
- Index change
- Landscape reorganization for SAP HANA scale-out systems
- Apply HINT to queries

Open SAP HANA capture and replay from the SAP HANA cockpit. In the [Performance Management](#) group, use the tiles [Capture Workload](#) and [Replay Workload](#).

For more information about SAP HANA capture and replay, see [Capturing and Replaying Workloads](#).

Related Information

[Capturing and Replaying Workloads \[page 410\]](#)

6.7.1.2.1 Capturing and Replaying Workloads

Capturing and replaying workloads from an SAP HANA database helps you evaluate potential impacts on performance or stability after a change in the hardware or software configuration.

The following sections provide an overview of SAP HANA capture and replay:

What is SAP HANA capture and replay?

This performance management tool allows you to capture the workload of a source system and to replay the captured workload on a target system without applications.

Moreover, you can use the tool to analyze the captured workload and the reports generated after replaying the workload. Comparing the performance between the source and target systems can help you find the root cause of performance differences.

What is a workload?

Workload in the context of SAP HANA can be described as a set of requests with common characteristics. For more information about workload in the context of SAP HANA, see *Workload in the Context of SAP HANA*.

In the context of SAP HANA capture and replay, workload can mean any change to the database via SQL statements that come from SAP HANA client interfaces such as JDBC, ODBC, or DBSL. The workload can be created by applications or clients (for example, SAP NetWeaver or Analytic).

How does SAP HANA capture and replay work?

The main steps involved in the capturing and replaying process are:

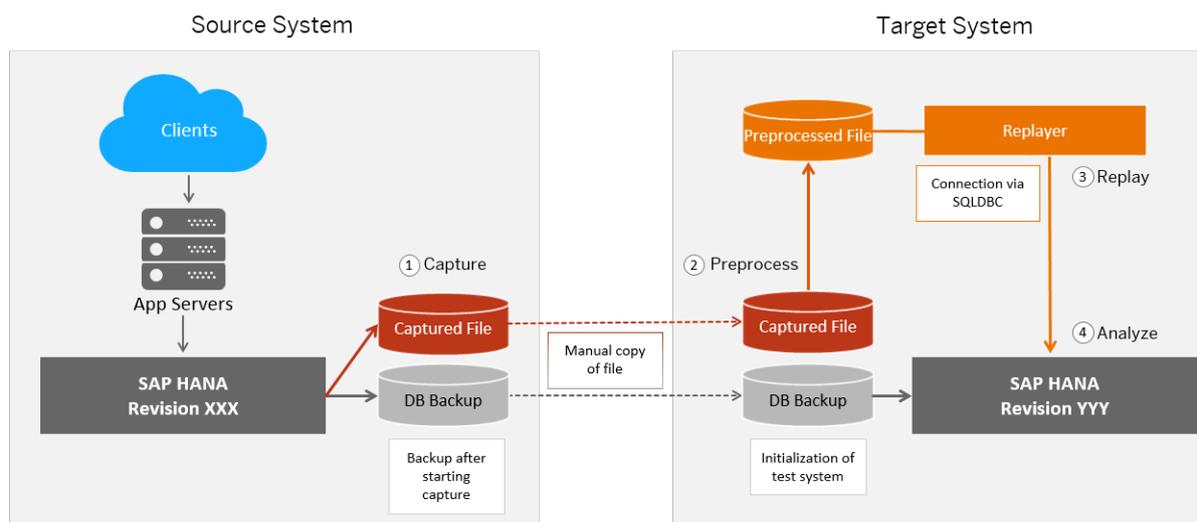
1. Capture
In this step the tool automatically collects the execution context information together with the incoming requests to the database. The captured file stores the start times of the SQL statements.
A full database backup is recommended after starting capturing to ensure that the source and target systems are in a consistent state.
2. Preprocess
In this step the tool reconstructs and optimizes the captured file to make it replayable on a target system. This process is a one-time operation and the stored preprocessed file can be replayed multiple times.
3. Replay
The replayer is a service on operating system level that needs to be started before replaying.
The tool replays the preprocessed file based on the SQL statement timestamp or on the transactional order. Together with the collected execution context it allows you to accurately simulate the database workload.
4. Analyze

For a final analysis, you can generate comparison reports displaying a capture-replay or a replay-replay comparison. You can analyze the statements based on results or on performance.

The following interactive graphic gives an overview of the main steps involved in the capturing and replaying process.

How does SAP HANA capture and replay work?

Hover over each numbered step for a description. Click the numbered steps for more information.



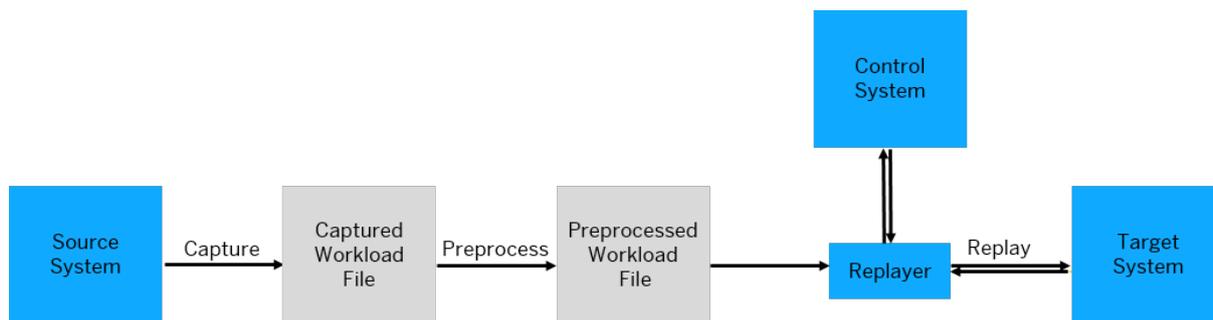
- [Capture a Workload \[page 416\]](#)
- [Preprocess a Captured Workload \[page 421\]](#)
- [Replay a Preprocessed Workload \[page 424\]](#)
- [Analyzing Comparison Reports \[page 430\]](#)

How to prepare your system landscape for SAP HANA capture and replay?

For the setup you can use 2 or 3 systems.

In a 2-system setup you need a source and a target system. In a 3-system setup, a control system should be added. This control system is the system running the cockpit and the system storing intermediate preprocessed or replay results. In a 2-system setup, the control system and the target system are shared. The advantage of a 3-system setup is that the replay results will be stored in the separate control system when recovering the target system.

The graphic below offers a visualization of a 3-system setup:



→ Recommendation

Recommendations when using SAP HANA capture and replay:

- Check the disk performance to ensure that there is sufficient bandwidth for capturing and preprocessing workloads without any performance bottlenecks. Otherwise, the instance that is running in capture mode can impact the source system.
- Check the available disk space in combination with the characteristics of the workload that should be captured. The required disk space is highly dependent on the type of workload being captured. Use the disk space that is dedicated to the database instance itself.
- One replayer service is sufficient to execute a replay successfully. For a better scalability and performance in large workload scenarios, multiple replayers can be used for all replaying purposes. When using multiple replayers distribute and divide all involved components (for example, target instance, control instance, one or more replayers) on different hosts and systems. Doing so will reflect the initial captured workload as realistic as possible and reduce the effect which the resource consumption of the components might have on a replay.
- Use a separate control and target instance for replaying workloads. If a replayed statement causes a crash, it will be displayed in the replay report. When you use one and the same control and target instance, the replay report entry causing a crash will not be successfully sent to the control instance.
- Use the Secure Store for saving passwords and authenticating users.

Which system privileges do you need?

You need the following system privileges:

- WORKLOAD CAPTURE ADMIN for capturing workloads
- WORKLOAD REPLAY ADMIN for preprocessing, replaying workloads, and viewing the load chart in the replay report
- WORKLOAD ANALYZE ADMIN for viewing the load chart in the replay report

Related Information

[SAP Note 2362820](#)

[Managing Performance in SAP HANA Cockpit \[page 409\]](#)

[Capturing a Workload \[page 415\]](#)

[Replaying a Workload \[page 419\]](#)

[Workload in the Context of SAP HANA \[page 621\]](#)

6.7.1.2.2 Key Performance Indicators for SAP HANA Capture and Replay (work in progress - tbd)

SAP HANA capture and replay allows you to select a range KPIs to ...tbd...

Comparison Report KPIs

KPI	Description
Elapsed Time	
Compile Time	
Wait Time	
Fetch Time	
CPU Time	
Lock Wait Duration	
Statement Dependency Wait Time	
Execution Delay	
Networking Sending Time	
Networking Receiving Time	
Execution Time	Query execution time inside the query layer.
Execution Close Time	Cursor fetch time inside query layer.
Execution Open Time	Cursor open time inside query layer.
Execution Fetch Time	Cursor fetch time inside query layer.

6.7.1.2.3 SAP HANA Capture and Replay Configuration Parameters

Several configuration parameters are available for SAP HANA Capture and Replay.

Parameter	<code>request_preload_size</code>
Value	
Default	
Description	<p>It is responsible for loading data before the replay for each captured connection. For example, if set to 1, each connection will load 1 request at the beginning of each replay.</p> <p>The requests will be loaded on demand. However, if the workload has a lot of idle time between requests, you can set a low value to avoid negative effects on the replay time. If you have high intensity workloads at the beginning of the replay, the overall replay time might be longer because only 1 request is loaded per connection.</p>

Parameter	<code>request_load_buffer_size</code>
Value	Number of requests
Default	
Description	<p>It is used to configure the buffer size used during replay.</p> <p>If you set a low value, it will take more time to replay. This is because the loading of requests into the buffer during replay will be triggered more often.</p>

Parameter	<code>replay_connection_packet_size</code>
Value	Value in MB
Default	for SQLDBC ~1 MB for CNR SQLDBS ~10MB
Description	<p>It is used for the SQLDBC connection.</p> <p>Each connection has a defined amount of memory assigned for each package it can send via SQLDBC. If the requests are too large to fit into the assigned package size, SQLDBC will split them into smaller request packages affecting the overall replay performance. This is because during replay, additional information is transmitted using SQLDBC. Lowering the value can significantly improve memory consumption, but make sure that the package size is big enough for the requests and any other additional information sent to the database during replay.</p>

6.7.1.2.4 SAP HANA Capture and Replay Filters

Some useful filters are available with which you can customize the amount and type of information displayed.

Filter	Description	Page	Tab
Application Name			
Application User Name			
Database User Name			
Client			

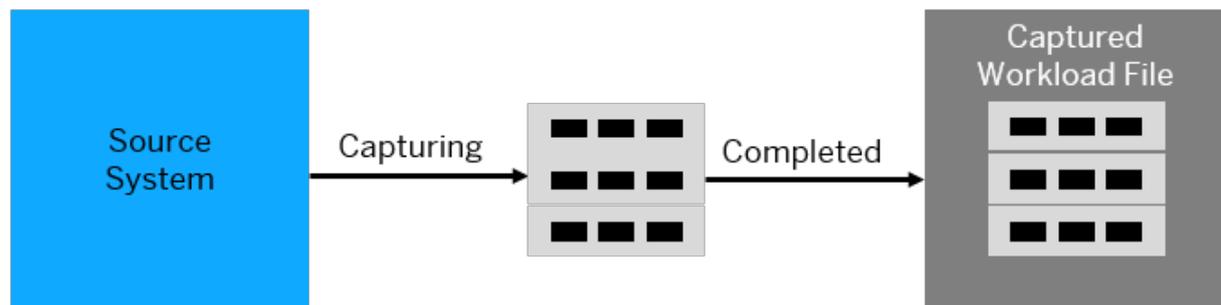
6.7.1.2.5 Capturing a Workload

You can capture the entire workload from a source system or only a part of this workload.

To capture the workload from a source system, use the *Capture Workload* tile in the *Performance Management* group.

After capturing the workload from a source system, a captured workload file will be available for the replay. This captured workload file contains multiple captured workload segment files. On a conceptual level, we will refer to the captured workload file using the shorter term *captured workload*.

The graphic below visualizes the concepts and the terminology:



Related Information

[Capture a Workload \[page 416\]](#)

[Capture Configuration Settings \[page 417\]](#)

6.71.2.5.1 Capture a Workload

You can capture and monitor the workload from a source system.

Prerequisites

You have a user with the WORKLOAD CAPTURE ADMIN system privilege.

Additionally you can add the following optional privileges:

- INIFILE ADMIN privilege allows you to see the previously used optional filters on the capture configuration page
- BACKUP OPERATOR privilege allows you to make a backup from the database

Procedure

1. On the *Overview* page, choose the *Capture Workload* tile.

The *Capture Management* page opens. If you already captured workload with SAP HANA capture and replay, you see the captured workload located on the current system.

2. To change INI configuration parameters, choose *Configure Capture* on the bottom right.

You can change the following parameters:

- *Capture Destination*
The captured workload file is stored by default in the \$SAP_RETRIEVAL_PATH/trace directory. Since the default trace directory generally resides in the same storage area with data and log volumes, capturing workloads may affect the performance across the entire system over time. Enter a different destination for the captured workload file to have a better distribution of the disk I/O between the data and log volumes, and the captured workload file.
- *Trace Buffer Size*
Change the size of the trace buffer that contains the captured workload file before it is written to disk.
- *Trace Segment Size*
After the maximum segment size has been reached, a new segment will be created if the capture is still ongoing.

3. To start configuring the new capture, choose *Start New Capture* on the *Capture Management* page on the bottom right.

On the *Configure New Capture* page it is mandatory to enter the name of the new capture.

You can customize other optional settings before you start the capture. For an overview of these settings, see [Capture Configuration Settings](#).

4. Choose [Start New Capture](#) on the bottom right.

The [Capture Monitor](#) opens displaying monitoring information such as duration, the number of captured statements, or disk space. You can stop the capture or you can let it run as long as you wish. If you didn't create a backup when starting the capture, you can also start a full backup from the [Capture Monitor](#) page.

i Note

The captured workload file is stored under the trace directory ((by default \$DIR_INSTANCE/<host name>/trace) with an *.cpt file extension.

You can return to the [Capture Management](#) page. After the capture is completed, the new captured workload has the status [Captured](#). By clicking the new captured workload, the [Capture Report](#) opens displaying information about the captured workload. You can continue analyzing the captured workload. For more information, see [Analyze a Captured Workload](#).

Related Information

[Capture Configuration Settings \[page 417\]](#)

[Analyze a Captured Workload \[page 418\]](#)

6.7.1.2.5.1.1 Capture Configuration Settings

You can customize several optional settings before you start capturing a workload.

The table provides an overview of the optional settings that can be customized.

Capture Configuration Settings

Setting	Description
Description	Enter a description of the capture for future reference. You can use this information as differentiator when capturing different scenarios in the same system.
Schedule	Schedule the capture by specifying the start and end time.
Overwrite Capture When Time Exceeds	Turn it on and enter a time to remove the captured workload segment files that are older than the specified time you entered. Only closed segments are deleted. The currently active captured workload segment file is not affected.

Setting	Description
Overwrite Capture When Disk Usage Exceeds	<p>Turn it on and select a ratio to remove the old captured workload segment files when the disk usage exceeds the specified percentage.</p> <p>Only closed segments are deleted. The currently active captured workload segment file is not affected.</p>
Collect Explain Plan	<p>Turn it on to collect the output of the EXPLAIN PLAN command for the captured statements.</p> <p>You can use this information for analysis after the replay.</p>
Collect Workload Details	<p>Turn it on to collect additional information for the instrumentation-based workload analyzer such as application source, involved threads, network statistics, or related objects.</p> <p>If this option is disabled, the captured workload file can still be viewed using the instrumentation-based workload analyzer, but less information will be available for the review.</p>
Create Full Backup	<p>Turn it on to automatically create a full database backup after starting the capture.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin-top: 10px;"> <p>→ Recommendation</p> <p>To ensure that the source system and the target system are in a consistent state for capture and replay, we recommend to perform a full database backup after starting the capture. A full database backup is required only for the first time, because incremental backups can be used once the system has been initialized for the first time. For more information, see <i>SAP HANA Backup and Recovery</i>.</p> </div>
Optional Filter	<p>Select additional filters to capture only desired aspects of the workload.</p> <p>Filters can include different aspects such as <i>Application Name</i>, <i>Database User Name</i>, <i>Statement Hash</i>.</p>

Related Information

[SAP HANA Database Backup and Recovery \[page 1229\]](#)

6.7.1.2.5.1.2 Analyze a Captured Workload

You can analyze the captured workload using the workload analyzer based on engine instrumentation.

Context

You can use the workload analyzer based on engine instrumentation to analyze the captured workload before starting the replay.

Procedure

1. On the *Capture Management* page, choose the captured workload that you want to analyze.
2. On the *Capture Report* page, choose *Analyze Workload* on the bottom right to open the workload analyzer based on engine instrumentation.

For more information about the workload analyzer based on engine instrumentation, see *Analyzing Workloads* and *Analyze Workloads Based on Engine Instrumentation*.

i Note

If you want to see the SQL statement parameters in the replay report, load the .cpt file after opening the workload analyzer based on engine instrumentation.

Related Information

[Analyzing Workloads \[page 435\]](#)

[Analyze Workloads Based on Engine Instrumentation \[page 438\]](#)

6.7.1.2.6 Replaying a Workload

You can replay the preprocessed workload based on the SQL statement timestamp or on the transactional order.

Replaying a workload implies that the captured statements are executed again.

→ Recommendation

Manually copy the captured workload files and the database backup from the source system to the control or target system.

You can replay all captured workloads as often as necessary.

→ Recommendation

When running consecutive replays, we recommend to restore the target system back to a consistent state after a replay and before running another replay. This is necessary because after replaying a workload on a system, any changes applied during that replay will remain active in the system.

❖ Example

Let's assume the captured workload file includes the statement `<INSERT INTO TABLE A VALUES (x).` During a replay, value `x` will be inserted into table A. At the end of the replay, table A contains value `x`. If you run another replay without resetting the system to its initial state, table A will contain duplicate values `x, x` at the end of the replay or the statement will fail (for example, in the case of unique constraint errors). As

x, x does not reflect the intended end result of the replay, it is recommended to restore the system after every replay when running multiple replays of the same captured workload.

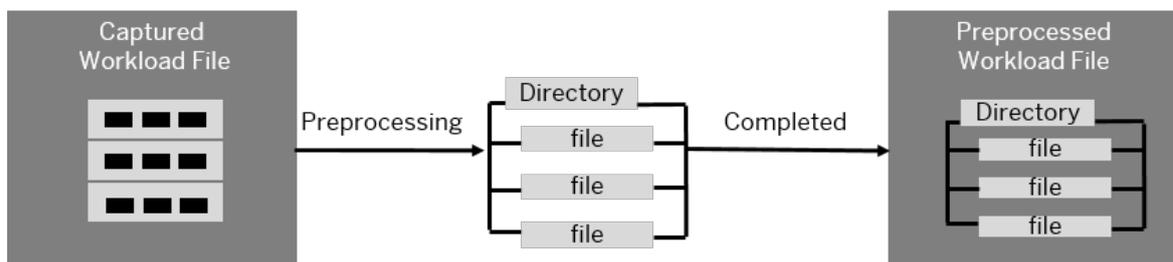
The following two steps are necessary before replaying the captured workload:

1. Preprocess a captured workload

The preprocessing step is required to optimize the captured workload file before replaying it. For more information, see *Preprocess a Captured Workload*.

While preprocessing the captured workload file, the captured statement segment files are stored in a directory. After the preprocessing is completed, the output is a preprocessed workload file containing the directory with multiple files. On a conceptual level, we will refer to the preprocessed workload file using the shorter term *preprocessed workload*.

The graphic below visualizes the concepts and the terminology:



2. Start the replayer

The replay process is performed by the replayer, which should be running before starting the replay. The replayer is a service on operating system level that reads SQL commands from the preprocessed workload file and executes them one-by-one in timestamp-based order. For more information, see *Start and Stop the Replayer*.

You can preprocess a captured workload and replay the preprocessed workload using the [Replay Workload](#) tile.

Related Information

[Preprocess a Captured Workload \[page 421\]](#)

[Start and Stop the Replayer \[page 422\]](#)

[Replay a Preprocessed Workload \[page 424\]](#)

[Replay Configuration Settings \[page 425\]](#)

[Generating Comparison Reports \[page 428\]](#)

[Generate a Replay-Replay Comparison Report \[page 429\]](#)

[Analyzing Comparison Reports \[page 430\]](#)

6.7.1.2.6.1 Preprocess a Captured Workload

The preprocessing step is necessary to optimize the captured workload file before replaying it.

Prerequisites

- You have a user with the WORKLOAD REPLAY ADMIN system privilege.
- You have captured workloads using the [Capture Workload](#) tile. For more information, see [Capture a Workload](#).
- Copy the captured workload file from the production system to the target system in the trace directory. If you use a control system that is different from the target system, copy the captured workload file to the control system.

→ Recommendation

We recommend to perform the preprocessing step in the target system or in a separate control system, not in the production system. The preprocessing may require significant computing power.

Procedure

1. On the [Overview](#) page, choose the [Replay Workload](#) tile.

The [Replay Management](#) page opens displaying an overview of the replay candidates located on the current system.

2. Check the captured workload that you want to preprocess.
3. To change INI configuration parameters, choose [Configure Replay](#) on the bottom right.

You can change the following parameters:

- [Preprocess Destination](#)
After the preprocessing is completed, the preprocessed workload file is stored by default in the \$SAP_RETRIEVAL_PATH/trace directory. Since the default trace directory generally resides in the same storage area with data and log volumes, preprocessing workloads may affect the performance across the entire system. Enter a different destination to have a better distribution of the disk I/O between the data and log volumes, and the preprocessed files.
- [Replay Report Elapsed Time Threshold](#)
The default value is set to 100 ms. When dealing with statements that display minor runtime differences, the tolerance ratio can lead to unexpected results after validating the delta. You can configure this field to avoid unintended and misleading classifications.

❖ Example

Executed statement A has a runtime of 10ms during capture, while executed statement A has a runtime of 12ms during replay. If the runtime in the target system is lower than the configured threshold value, the statement will no longer be listed as *slower* or *faster*, but as *comparable* in the replay report.

4. Click [Start Preprocessing](#) on the bottom right.

Related Information

[Capture a Workload \[page 416\]](#)

6.7.1.2.6.2 Start and Stop the Replayer

The replayer is a service on operating system level that should be running before starting the replay.

Prerequisites

- You have a user with the WORKLOAD REPLAY ADMIN system privilege to control the replayer. Store the logon credentials in the secure store. For more information, see *Secure User Store (hdbuserstore)* in the *SAP HANA Security Guide*.
- When using multiple replayers distribute and divide all involved components (for example, target instance, control instance, one or more replayers) on different hosts and systems.

i Note

The replayer is not a part of the SAP HANA database services that are running as daemon process. You must start and stop it yourself.

Procedure

1. Configure a hdbuserstore entry to authenticate the replayer with the database using the following command on the operating system on which your SAP HANA database is installed:

```
hdbuserstore SET <key name> <host name@tenant database name> <user name>
<password>
```
2. Start the replayer using the following command on the Linux command line of the system that you want to start the replayer on:

```
hdbwltreplayer -controlhost <controlHost> -controlinstnum <controlInstanceNumber>
-controladminkey <userName,secureStoreKey> -controldbname <controlDatabaseName>
-port <listenPortNumber>
```

i Note

Running the command on the target system does not trigger the replay, it only starts the replayer.

The parameters `controlhost`, `controlinstnum`, `controladminkey`, and `controldbname` indicate the location of the control system.

Parameter Description

Parameter	Description
<code>controlhost</code>	Specifies the database host name of the control or target system (without a <code>sqlport</code>).
<code>controlinstnum</code>	Specifies the database instance number.
<code>controladminkey</code>	Specifies the user name and secure store key of the control management connection separated by a comma.
<code>controldbname</code>	Specifies the database name. When connected to a tenant, the tenant name should be used. When connected to a system database, the system database should be used as control database name.
<code>port</code>	Specifies the discretionary port number for internal communication.

To start multiple instances of the replayer in parallel, define a specific port for each instance. The setup of the second instance fails when two instances run on the same port.

When running replays with a longer duration, you can add `&` at the end of the command line. This starts the process in the background and you can close the terminal connection immediately.

❁ Example

```
hdbwltreplayer -controlhost <controlHost> -controlinstnum  
<controlInstanceNumber> -controladminkey <userName,secureStoreKey> -  
controldbname <controlDatabaseName> -port <listenPortNumber> &
```

i Note

If you want to use SSL encryption with the replayer, navigate to the `wltreplayer.ini` file on the OS and edit or add the section `[replay_client]` as follows:

- Add the parameter `enable_target_ssl_connection = [true|false]` to enable SSL connections between the target system and the replayer
- Add the parameter `enable_control_ssl_connection = [true|false]` to enable SSL connections between the control system and the replayer

Only fully qualified domain names can be used for `<controlHost>` when starting the replayer.

3. If the console is still open, use `Ctrl+C` to stop the replayer. Alternatively, identify the OS process ID of the running replayer and shut it down using `kill<pid>`.

6.7.1.2.6.3 Replay a Preprocessed Workload

You can replay all preprocessed workloads as often as necessary.

Prerequisites

- You have a user with the WORKLOAD REPLAY ADMIN and WORKLOAD ANALYZE ADMIN system privileges.
- The target system meets the same security and privacy prerequisites as the source system. Since the target system processes the same data as the source system, it should meet an appropriate security level depending on data criticality.

→ Recommendation

Do not allow unnecessary network connections to the target system. Users registered on the source system could access the target system after the replay is completed.

- You have preprocessed the captured workloads using the [Replay Workload](#) tile. For more information, see [Preprocess a Captured Workload](#).
- The replayer is running. For more information, see [Start and Stop the Replayer](#).

Procedure

1. On the Overview page, choose the [Replay Workload](#) tile.

The [Replay Management](#) page opens displaying the captured workload.

2. To change INI configuration parameters, choose [Configure Replay](#) on the bottom right.

You can change the following parameters:

- [Preprocess Destination](#)
After the preprocessing is completed, the preprocessed workload file is stored by default in the \$SAP_RETRIEVAL_PATH/trace directory. Since the default trace directory generally resides in the same storage area with data and log volumes, preprocessing workloads may affect the performance across the entire system. Enter a different destination to have a better distribution of the disk I/O between the data and log volumes, and the preprocessed files.
- [Replay Report Elapsed Time Threshold](#)
The default value is set to 100 ms. When dealing with statements that display minor runtime differences, the tolerance ratio can lead to unexpected results after validating the delta. You can configure this field to avoid unintended and misleading classifications.

❖ Example

Executed statement A has a runtime of 10ms during capture, while executed statement A has a runtime of 12ms during replay. If the runtime in the target system is lower than the configured threshold value, the statement will no longer be listed as *slower* or *faster*, but as *comparable* in the replay report.

3. Choose a replay candidate with the status *Completed* to start configuring it for the replay.

The *Replay Configuration* page opens allowing you to configure various mandatory and optional settings. For more information about each setting, see *Replay Configuration Settings*.

i Note

If a database backup is available, restore the database before starting the replay in the target system. For more information, see *SAP HANA Backup and Recovery*. When running a replay on a target system that has been restored using a backup taken automatically during the capture process, activate the *Synchronize Replay with Backup*.

If no or only outdated database backups are available, you can still restore the database or manually export parts of the data before starting the replay in the target system. When running a replay on a target system that was restored using old backups or contains only smaller manual exports of data, deactivate the *Synchronize Replay with Backup* option.

4. After configuring the replay, choose *Review* to view the replay configuration.
5. To start the replay, choose *Confirm*.

The *Replay Management* opens displaying in the *Replay List* tab the workloads that are being replayed. You can start multiple replays in parallel. For more information on each replay, open the *Details* link in the message field. If you don't want to see this information on your *Replay List* anymore, use the button on the top right to close it. To stop a replay when it is in progress, choose *Stop* in the *Replay Status* column. The status of the replay changes then to *Stopped*.

If you want to import a replay from a different system, use the *Import Replay* button on the bottom right.

To access the *Replay Monitor*, choose the running replay. The monitoring view provides information such as duration, number of statements, size, and other details about the replay in progress. You can navigate away from the monitoring view using the arrow on the top right and can return anytime.

If you have already replayed preprocessed workloads, you can generate comparison reports for further analysis. For more information, see *Generating Comparison Reports*.

Related Information

[Preprocess a Captured Workload \[page 421\]](#)

[SAP HANA Database Backup and Recovery \[page 1229\]](#)

[Start and Stop the Replayer \[page 422\]](#)

[Replay Configuration Settings \[page 425\]](#)

[Generating Comparison Reports \[page 428\]](#)

6.7.1.2.6.3.1 Replay Configuration Settings

The *Replay Configuration* allows you to set various parameters to best suit how the replay is to be processed.

The *General Information* page allows you to customize the following mandatory and optional settings:

General Replay Information

Setting	Replay Name
Type	Mandatory
Description	Enter a name for the replay. By default this field has the same name as the initial captured file.

Setting	Description
Type	Optional
Description	Enter a description for the replay for your future reference. This information can be used when changing settings for different replays.

Target System Information

Setting	Host
Type	Mandatory
Description	Enter the target host name (for example, ld2536) where the capture will be replayed. This field defines the target host name of the system on which the replay should run.

Setting	Instance Number
Type	Mandatory
Description	Enter the target instance number (for example, 77) where the capture will be replayed. This field defines the target instance number of the database on which the replay should run.

Setting	Container
Type	Mandatory
Description	Select between a Single Container or Multiple Containers . Specify on which tenant the replay should run.

Replayer Options

Setting	User Name
Type	Mandatory
Description	Enter the database user who has the WORKLOAD REPLAY ADMIN privilege and will be used for the final preparation steps in the target instance.

Setting Request Rate

Type Optional

Description Modify the rate at which the statements are replayed.

You can decrease the wait time between statements during replay. For example, statement B starts 1second after statement A has been triggered. When setting the request rate from *1x* to *2x*, this difference will only be 0.5 seconds.

Setting Synchronize Replay with Backup

Type Optional

Description This option allows you to synchronize the replay with an existing database backup.

The option is turned on by default allowing the replayer to compare each statement with the database backup. This option makes it possible to check if there are no duplicate inserts and if the backup and replay are aligned. A backup is required for this option to work correctly.

If the option is turned off, the replayer will replay statements, even if no backup is present. This is important for scenarios in which you use only single tables, or smaller data exports, which are not considered a complete backup.

Setting Collect Explain Plan

Type Optional

Description Collect the output of the EXPLAIN PLAN command for captured statements.

You can use this information for comparison after the replay.

Setting Transactional Replay

Type Optional

Description This option enables guaranteed transactional consistency during a replay.

i Note

Enabling this option may cause overhead to query runtime as transactional consistency needs to be checked constantly.

Optional Filter

Setting Optional Filter

Type Optional

Setting **Optional Filter**

Description Use the *Optional Filter* to selectively replay only desired aspects of workloads.
Filters can include different aspects such as *Application Name*, *Database User Name*, *Statement Hash*.

The *Replay Information* page allows you to customize the following mandatory settings:

Setting **Replayer List**

Type Mandatory

Description Select a running *Replayer* that will be used to connect to the target system and facilitate the replay.

Setting **User Authentication**

Type Mandatory

Description Enter the password or the securestore key for the database users captured in the source system. For a realistic replay, all users that are part of the workload, which has been chosen to be replayed, must be authenticated. For more information, see *Secure User Store (hdbuserstore)* in the *SAP HANA Security Guide*.

To reset the password for the database users captured in the source system, select the users, and then choose *Reset Password*. This can be helpful when you don't know the actual password of each user. On the *Reset Password* window, set a new password for all selected users, and choose *Confirm*. All selected user passwords in the defined target system will be changed as defined in this step.

6.7.1.2.6.3.2 Generating Comparison Reports

You can generate comparison reports after successfully replaying a captured workload.

You can generate comparison reports displaying a capture-replay or a replay-replay comparison.

On both types of reports you can perform the following actions:

- Download detailed information in JSON format by choosing the save button at the top of the statement detail table. You can choose the category and the number of statements that you want to select.
- Export replay reports to store them outside the database. Choose the arrow on top right to export the replay results.
- Import exported replay reports using the *Import Replay* button at the top of the *Replay List* tab. For information on the security implications and configuration steps needed for importing replay reports, see SAP Note 2109565.

i Note

Storing the replay results outside the database can be useful when the target and control systems are the same. In such a setup, the previous replay results of the control system could be overwritten after recovering the target system from the database backup.

Capture-Replay Comparison Report

You can open a comparison report by choosing a replay from the [Replay List](#). When opening a comparison report directly from the [Replay List](#), the report shown always compares values from the original captured workload with values from the replay.

Replay-Replay Comparison Report

When using the [Compare Replays](#) button on the bottom right, the report shown always compares different replays with each other based on the same initial captured workload.

For more information about generating a report based on a replay-replay comparison, see [Generate a Replay-Replay Comparison Report](#).

Related Information

[Generate a Replay-Replay Comparison Report \[page 429\]](#)

[SAP Note 2109565](#) 

6.7.1.2.6.3.2.1 Generate a Replay-Replay Comparison Report

You can compare two or more replayed workloads with each other based on the same initial captured workload.

Prerequisites

- You have a user with the WORKLOAD REPLAY ADMIN and WORKLOAD ANALYZE ADMIN system privileges.
- You have replayed preprocessed workloads using the [Replay Workload](#) tile. For more information, see [Replay a Preprocessed Workload](#).

Context

You can start the comparison of the replayed workloads from the [Replay List](#) in the [Replay Management](#).

For more information about the [Replay List](#), see [Replay a Preprocessed Workload](#).

Procedure

1. On the *Replay List* tab, click *Compare Replays* on the bottom right.

The *Select Baseline Replay* dialog opens, allowing you to select the replayed workload that you want to compare. Use the *Target SID* information to distinguish between the replays.

2. Select one entry from the displayed list and click *Close*.

The *Select Target Replay* dialog opens allowing you to select the replayed workload that you want to compare with the previously selected workload. The list displays replayed workloads based on the same initial captured workload.

3. Select one or more entries from the displayed list and click *Compare Replays* on the bottom right.

The *Comparison Report* opens displaying a comparison of the selected replayed workloads.

Related Information

[Replay a Preprocessed Workload \[page 424\]](#)

[Analyzing Comparison Reports \[page 430\]](#)

6.7.1.2.6.3.3 Analyzing Comparison Reports

You can use comparison reports to analyze the completed replay.

The information is displayed on four tabs:

Overview

The *Overview* tab displays an overall comparison of the SQL statements involved in the capturing and replaying process in the following blocks:

Overview Information

Block Name	Description
Result Comparison	In a result-based comparison you get an overview of the statements with identical or different results. Click the block to open directly the <i>Result Comparison</i> tab.
Performance Comparison	In a performance-based comparison you get an overview of the statements based on a comparison of run-times. Click the block to open directly the <i>Performance Comparison</i> tab.

Block Name	Description
Different Statements	Displays the top SQL statements that have different results from the selected baseline in descending order. You can click each row to open the Execution Detail page for the selected SQL statement. Use the drop down arrow to filter the statements by time or by the number of records that have different results.
Slower Statements	Displays the top SQL statements that have a different performance ordered by the difference in execution time. You can click each row to open the Execution Detail page for the selected SQL statement. To view KPI details for each statement, you can click the icon on the right.
Verification Skipped	Displays the distribution of reasons for statements with skipped result comparison.
Replay Failed Statements	Displays the distribution of reasons for the statements, which failed during replay. Use the drop down arrow to filter the statements by time or error code.
Capture Information	Displays information on the capture system, capture options, and the properties of the capture file.
Replay Information	Displays information on the replay system and the replay options. If the comparison was made between two replays, the information is displayed in a Baseline Replay Information block and in a Target Replay Information block.

To generate a workload replay report in PDF format, choose [Generate PDF](#) on the bottom right.

Load

This tab includes load charts comparing both the captured and the replayed workloads after a capture-replay comparison, or the baseline and the target workloads after a replay-replay comparison. The KPIs can be toggled independently for both the capture and replay aspects, making it easier to compare them with each other. Additional KPIs can be added using the [Show More KPIs](#) button on top right of the load chart.

To generate a workload replay report in PDF format, choose [Generate PDF](#) on the bottom right.

Performance Comparison

The performance-based comparison provides an overview of statements compared by runtime.

Based on the selected tolerance ration, the statements are classified as [Comparable](#) when they have a similar runtime within the defined tolerance ratio, [Faster](#), [Slower](#), or [Failed](#). For more information about the tolerance ratio, see [Replay Report Elapsed Time Threshold](#) in *Preprocess a Captured Workload*. For more information on how the statements are classified, see *Statement Classification*.

The list below the graphic displays a detailed view of the individual statements. You can use the fields and buttons on top to:

- View the defined tolerance ratio.
- Search for a specific statement.
- Display more statements.
- Refresh the displayed data.
- Change the sorting criteria.
- Export detailed information in JSON format by clicking the save button.
- Add more columns for analysis.

To view a summary of how the execution time was spent, select the statistics symbol at the end of each statement line.

To open the execution details for a specific statement, select the statement from the list. On the [Execution Details](#) page, you can:

- Create a list of further statements, by choosing [Show More Statements](#) on the top right.
- Search the list of individual executions.
- View any parameters that might have been part of the query. To display the parameters, the associated captured workload file must be loaded in the instrumentation-based workload analyzer.
- View the runtime [KPIs](#) for that execution. For more information on each key performance indicator, see [Key Performance Indicators for SAP HANA Capture and Replay](#).
- Display the explain plan comparison for that statement by clicking on it.

Result Comparison

The result-based comparison provides an overview of statements with, for example, identical or different results.

The result-based replay report also includes a classification of statement types based on the content of those statements being either deterministic or non-deterministic. Deterministic statements should always deliver the same results during a replay. Non-deterministic statements are expected to deliver different results (for example, because they don't contain an explicit sorting of results).

The list below the graphic displays a detailed view on the individual statements. You can use the fields and buttons on top to:

- Search for a specific statement.
- Display more statements.
- Refresh the displayed data.
- Change the sorting criteria.
- Export detailed information in JSON format by clicking the save button.

To open the execution details for a specific statement, select the statement from the list.

You can use the detailed execution level of both report types to compare the EXPLAIN PLAN results between the initial captured and replayed workloads, or between the baseline and target workloads. Comparing the plans can provide guidance and pointers for further statement-level investigation. This is only possible if the [Collect Explain Plan](#) setting was activated for both the capture and the replay during the configuration steps. For more information about this setting, see [Capture Configuration Settings](#) and [Replay Configuration Settings](#).

Related Information

[Replay Configuration Settings \[page 425\]](#)

[Capture Configuration Settings \[page 417\]](#)

[Preprocess a Captured Workload \[page 421\]](#)

[Statement Classification \(work in progress - tbd\) \[page 433\]](#)

[Key Performance Indicators for SAP HANA Capture and Replay \(work in progress - tbd\) \[page 413\]](#)

6.71.2.6.3.3.1 Statement Classification (work in progress - tbd)

In the performance-based comparison the statements are classified as slower, faster, failed or comparable.

The classification of the statements in the performance-based comparison is based on the interplay between the values of the tolerance ratio, target elapsed time, and source elapsed time.

Comparable Statements

For comparable statements the value is calculated as follows:

```
Comparable=(Tolerance Ratio*-1)<(Target Elapsed Time-Source Elapsed Time)/  
(Source Elapsed Time)*100<(Tolerance Ratio)
```

❖ Example

With a tolerance ratio of 10%, a target elapsed time of 8.117 ms, and a source elapsed time of 8.458 ms:

```
(8.117-8.458)/8.458*100= -4.03
```

The statement is executed 4.03% faster on the target (replay) side compared to the source (capture) side. However, the statement is classified as comparable because the tolerance ratio is defined for ±10%.

Alternatively, the classification relies on the comparison between the target elapsed time and the replay report threshold:

```
Comparable=Target Elapsed Time<Replay Report Threshold
```

❖ Example

With a tolerance ratio of 10%, a target elapsed time of 1.400 ms, a source elapsed time of 1.100 ms, and a replay report threshold of 1500 µs (1.500 ms):

```
1.400<1.500=true
```

The statement is classified as comparable because the replay report threshold is higher than the target elapsed time.

Faster Statements

For faster statements the value is calculated as follows:

```
Faster=(Target Elapsed Time-Source Elapsed Time)<0  
and  
Faster=(Target Elapsed Time-Source Elapsed Time)/(Source Elapsed  
Time)*100<(Tolerance Ratio*-1)  
and  
Faster=Target Elapsed Time>Replay Report Threshold
```

❖ Example

With a tolerance ratio of 10%, a target elapsed time of 6.379 ms, a source elapsed time of 9.390 ms, and a replay report threshold of 700 μ s (0.700 ms):

$$(6.379-9.390)/9.390*100= -32.07$$

The statement is executed 32.07% faster on the target (replay) side compared to the source (capture) side. The statement is classified as faster because the tolerance ratio is lower and the replay report threshold is lower than the target elapsed time.

Slower Statements

For slower statements the value is calculated as follows:

```
Slower=(Target Elapsed Time-Source Elapsed Time)>0  
and  
Slower=(Target Elapsed Time-Source Elapsed Time)/(Source Elapsed  
Time)*100>Tolerance Ratio  
and  
Slower=Target Elapsed Time>Replay Report Threshold
```

❖ Example

With a tolerance ratio of 10%, a target elapsed time of 1.253 ms, a source elapsed time of 0.991 ms, and a replay report threshold of 700 μ s (0.700 ms):

$$(1.253-0.991)/0.991*100= 26.44$$

The statement is executed 26.44% slower on the target (replay) side compared to the source (capture) side. The statement is classified as slower because the tolerance ratio is lower and the replay report threshold is too low.

6.7.1.3 Analyzing Performance in SAP HANA Cockpit

You can analyze the performance of the database using the SAP HANA cockpit.

You can use the following tools to analyze fine-grained aspects of database performance in the SAP HANA cockpit:

- Use the *Workload Analyzer* to analyze the database performance through data from thread samples. You can also use it to analyze the workload captured with the capture and replay tool, or any other workload occurring in a system.
- Use the *SQL Analyzer* to understand performance issues of a query execution and other query execution aspects of the SAP HANA database.

Related Information

[Analyzing Workloads \[page 435\]](#)

[Analyzing Statement Performance \[page 440\]](#)

6.7.1.3.1 Analyzing Workloads

Analyzing workloads from an SAP HANA database with the workload analyzer can help you identify the root cause of performance issues.

The following sections provide an overview of the workload analyzer tool:

What is the workload analyzer tool?

The workload analyzer is a tool that allows you to analyze the workload captured with the capture and replay tool, or any other workload occurring in a system.

The workload analyzer has two versions:

- The workload analyzer based on thread samples
This version uses thread samples data to analyze the performance.
- The workload analyzer based on engine instrumentation
This version of the tool performs the analysis using the captured data containing all the statistics of the query execution. Furthermore, it allows a deeper analysis with a full set of execution details.
For more information on how to capture the workload of a system with the capture and replay tool, see *Capture a Workload*.

What is the workload analyzer based on thread samples?

The workload analyzer based on thread samples is a solution for analyzing database performance using thread samples.

The workload analyzer based on thread samples provides a workload analysis using different KPIs, and it offers the following information sets:

- On the upper part of the screen, the chart displays the system resource usage. The chart displays both a real-time and a historical analysis. The information displayed on the grey background represents the historical analysis of the workload. Both analysis types are based on the sampling data. However, the historical analysis contains only aggregated data.
- On the lower part of the screen, the main analysis chart offers the following three sections:
 - *Background Jobs*
This section displays in the main chart information on the job progress. The miniature chart shows the system load data within the time range specified in the upper load chart.
 - *SQL Statements*
This section displays the timeline view consisting of two charts: a miniature chart showing the system load data during the specified time range, and a timeline chart. You can expand the application name by clicking the arrow next to it to get a list of statements being executed by the selected application.
 - *Threads*
In this section the main stacked area chart displays a more detailed visualization of the chart on the upper part of the screen based on a selected dimension (for example, thread type) over a given period of time. The bar charts located at the bottom next to it display the top five statements consuming most of the threads during the given timeframe. Clicking a specific statement hash opens a dialog with detailed statement information.

For more information on how to use the workload analyzer based on thread samples, see *Analyze Workloads Based on Thread Samples*.

What is the workload analyzer based on engine instrumentation?

The workload analyzer based on engine instrumentation is a solution for analyzing database performances based on the workload captured with the capture and replay tool.

The workload analyzer based on engine instrumentation offers you two analysis types:

- On the upper part of the screen, the chart displays the system resource usage.
- On the lower part of the screen, the timeline analysis based on an application and statement level hierarchy enables you to evaluate anomalies in the timeline and identify potentially problematic statements.

In contrast to the workload analyzer based on thread samples, using this tool requires to trace the workload before analyzing the performance. Tracing all the workload by default is not recommended because it introduces an overhead to system performance.

For more information on how to use the workload analyzer based on engine instrumentation, see *Analyze Workloads Based on Engine Instrumentation*.

Why should you use the workload analyzer tool?

The workload analyzer gives you an overview of the system's health at a glance. Moreover, the tool helps you identify the root cause of performance issues either by a real-time analysis or by reviewing historical data.

How to access the workload analyzer tool?

You can access both versions of the workload analyzer from the SAP HANA cockpit as follows:

- To access the workload analyzer based on thread samples you have the following possibilities:
 - Open the [Analyze Workload Based on Thread Samples](#) tile on the SAP HANA cockpit.
 - or
 - Open the [Performance Monitor](#) page by clicking in the [Monitoring and Administration](#) tile group [Show All](#) on the [Memory Usage](#), [CPU Usage](#), or [Disk Usage](#) tile. On the [Performance Monitor](#) page open the expand button next to the name of the page and select [Workload Analysis](#).
- To access the workload analyzer based on engine instrumentation open the [Analyze Workload Based on Engine Instrumentation](#) tile in the [Performance Management](#) tile group.

Related Information

[Capture a Workload \[page 416\]](#)

[Analyze Workloads Based on Thread Samples \[page 437\]](#)

[Analyze Workloads Based on Engine Instrumentation \[page 438\]](#)

6.7.1.3.1.1 Analyze Workloads Based on Thread Samples

You can analyze database performance using the workload analyzer based on thread samples.

Prerequisites

You have the system privileges `CATALOG READ` and `INIFILE ADMIN`.

Procedure

1. On the [Overview](#) page, open the workload analyzer either using the [Analyze Workload Based on Thread Samples](#) tile or the [Performance Monitor](#) page. For more information, see the [How to access the workload analyzer tool](#) section in [Analyzing Workloads](#).

The workload analyzer opens displaying a chart on system resource usage on the upper part of the screen and a more detailed visualization distributed in three sections: *Background Jobs*, *SQL Statements*, *Threads*.

2. Analyze the displayed charts using the following features:
 - a. Customize the information displayed on the load graph on the upper part of the screen by selecting the wished KPIs, selecting the refresh period and the duration, or using the navigation buttons on the top right to set the desired time frame. The selected KPIs appear in the legend area on the left-side of the chart. For a list of all available KPIs, see *Key Performance Indicators*. Furthermore, you can import and export datasets in order to store the data in an application and to analyze it in another system.
 - b. On the lower part of the screen in the *SQL Statements* section click *edit* to select the dimension group which will be used for the drill-down analysis. In the *Dimensions* dialog specify the dimension group considering the application name and the statement hash. If there are more than two dimensions specified as a group, you can drill-down into the next level by double-clicking the label in the legend area. A list of the services where the statements were executed by the specified application will be displayed.
 - c. On the lower part of the screen in the *Threads* section you can specify multiple filters (for example, statement hash, thread type, or application source) in order to analyze only the data you are interested in.

Related Information

[Analyzing Workloads \[page 435\]](#)

[Key Performance Indicators \[page 249\]](#)

6.7.1.3.1.2 Analyze Workloads Based on Engine Instrumentation

You can analyze database performance using the workload analyzer based on engine instrumentation.

Prerequisites

You have the system privileges `WORKLOAD ANALYZE ADMIN`.

Context

The *Analyze Workload Based on Engine Instrumentation* tile provides information on the number traced files, as well as on the number of loaded files. In order to analyze the traced workload data, the file has to be loaded into the database.

Procedure

1. On the Overview page, choose the [Analyze Workload](#) tile in the [Performance Management](#) tile group.

The [Trace List](#) page opens displaying workloads with the status [Loaded](#) or [Traced](#). Only workloads with the status [Loaded](#) can be analyzed.

On this page, you can use the search field to search for further captured files.

2. Optional: To analyze performance issues that are reproduced in the production system, trace the workload data by clicking [Start Trace](#) on the bottom right.

The [Configure New Trace](#) page opens allowing you to provide the needed information for tracing. Choose if you want the trace to be overwritten in case it exceeds a pre-defined time or disk usage. Once you have entered the information and chosen filters, click on [Start Trace](#).

Optional: To schedule tracing a workload in the future, turn the [Schedule](#) option on and define start and end times. Click on [Register Trace](#).

Once the tracing started the [Trace List](#) page will be displayed again.

3. Optional: To check the tracing details click on the file with the status [Tracing](#).

The [Trace Monitor](#) page opens displaying overall tracing information such as tracing time, number of collected statements, trace size. To stop tracing, click on the [Stop Trace](#) button.

4. Optional: To obtain information on a [Traced](#) workload, select the desired workload.

The [Trace Information](#) page opens with information on the [Traced](#) workload.

5. Optional: To load the trace, open the [Trace Information](#) page by selecting the desired [Traced](#) workload. Choose [Load Trace](#) on the footer bar.

The [Trace List](#) page opens with the status of the desired trace updated to [Loaded](#).

6. Click on any workload with the status [Loaded](#) on the [Trace List](#) page to open the analysis view.

The [Workload Analysis](#) page opens, displaying two analysis charts and an information table.

The [Overview](#) tab displays three bar charts that list top five entries for each measurement. You can specify the measurement for each bar.

Optional: Customize the tabs displayed on the lower part of the screen by clicking [Edit](#) and selecting the desired dimensions.

7. Optional: Customize the information displayed on the load graph on the upper part of the screen by selecting the host and services or selecting the wished KPIs. The selected KPIs appear in the legend area on the left-side of the chart. For a list of all available KPIs, see [Key Performance Indicators](#). Moreover, you can set specific filters (for example, statement hash, thread type, or application source) in order to analyze only the data you are interested in.

8. Optional: To analyze plans and jobs of an SQL statement in detail, go to the [Workload Analysis](#) page and switch to the [SQL Statements](#) tab on the lower chart. On the chart, select the workload you want to analyze. Then click on the magnifying glass icon that appears at the top of the bar.

The [Plans and Jobs](#) tab opens displaying plan data in green and job data in orange. The corresponding request is highlighted in the statement table below, which provides additional information.

9. Optional: To perform a hierarchical analysis of data, go to the [Workload Analysis](#) page and switch to the [SQL Statements](#) tab on the lower chart. Click [Edit](#) next to the dimension group tabs and add Request ID.

The [Request ID](#) tab is added to the dimension groups.

To look up which request ID or statement hash belongs to which application, select the desired *Application Name* value. To return to the top-level hierarchy, click on the *Application Name* again.

10. Optional: In the timeline view on the lower part of the screen, get more information about the statement by clicking on the statement hash.

Related Information

[Capture a Workload \[page 416\]](#)

[Key Performance Indicators \[page 249\]](#)

6.7.1.3.2 Analyzing Statement Performance

Analyzing statement performance helps you understand performance issues of a query execution and other query execution aspects of the SAP HANA database.

The following sections provide an overview of the SQL analyzer tool:

What is the SQL analyzer tool?

The SQL analyzer is a query performance analysis tool of SAP HANA. The tool can be used to view detailed information on each query execution and can help you evaluate potential bottlenecks and optimizations for these queries.

How can you access the SQL analyzer?

You can open the SQL analyzer from the SAP HANA cockpit or from the SAP Web IDE for SAP HANA.

- From the SAP HANA cockpit there are six ways to open the SQL Analyzer:
 - Using the *Monitor expensive statements* link from the *Monitoring* link list.
On the *Overview* page, open the *Monitor expensive statements* link. The *Expensive Statements Trace* page opens, allowing you to identify which SQL statements require a significant amount of time and resources. Each statement string is provided with a *More* link, which opens the *Full SQL Statement* dialog. Click *Open in SQL Analyzer* to open the selected query with the SQL analyzer tool.
 - Using the *Open SQL plan cache* link from the *Monitoring* link list.
On the *Overview* page, open the *Open SQL plan cache* link. The *SQL Plan Cache* page opens, allowing you to manage registered statement hints. Each statement string is provided with a *More* link, which opens the *Full SQL Statement* dialog. Click *Open in SQL Analyzer* to open the selected query with the SQL analyzer tool.
 - Using the *Manage statement hints* link from the *DB Administration* link list.
On the *Overview* page, open the *Manage statement hints* link. The *Statement Hints* page opens, allowing you to manage registered statement hints. Each statement string is provided with a *More* link,

which opens the *Full SQL Statement* dialog. Click *Open in SQL Analyzer* to open the selected query with the SQL analyzer tool.

- Using the *Plan Trace* link from the *Alerting & Diagnostics* link list.
On the *Overview* page, open the *Plan Trace* link. The *Plan Trace* page opens. To view a set of statistics for each SQL statement collected in a given time frame, click the *Refresh* button. Each statement string is provided with a *More* link, which opens the *Full SQL Statement* dialog. Click *Open in SQL Analyzer* to open the selected query with the SQL analyzer tool.
- Using the *Execute SQL* link from the *Database Explorer* link list.
On the *Overview* page, open the *Execute SQL* link. The SAP HANA database explorer opens. In the menu of the *Run* button, click *Analyze SQL* to open the SQL analyzer.
- From the SAP Web IDE for SAP HANA you can open the SQL analyzer in a SQL Console by clicking *Analyze SQL* in the menu of the *Run* button.

Which views are supported by the SQL analyzer?

The following views are supported by the SQL analyzer:

SQL Analyzer Views

View	Description
Overview	This view provides an overview of the query execution including metadata for the analysis.
Plan Graph	This view provides graphical guidance to help you understand and analyze the execution plan of a SQL statement. In case of SQLScript, the SQLScript definition is also displayed.
Timeline	This view provides a complete overview of the execution plan based on the visualization of sequential time-stamps.
Tables in Use	This view provides a list of tables used during query execution and includes further details on tables, which can be used for further analysis. It can be used to understand which tables are needed to fulfill a given SQL statement execution.
Operators	This view provides a list of operators used during query execution and includes additional details about each operator, which can be used for analysis.
Statement Statistics	This view is only displayed in the case of a SQLScript. This view provides details on the statistics of the executed statements in the SQLScript.
Table Accesses	This view provides details on the table accesses performed during the processing of a SQL statement, which can be used for analysis.

Related Information

[Analyze Statement Performance \[page 442\]](#)

[Monitor and Analyze Statements with Plan Trace \[page 445\]](#)

6.7.1.3.2.1 Analyze Statement Performance

The SQL analyzer is used to analyze statement execution performance.

Context

From the SAP HANA cockpit the SQL analyzer can be opened using the [Monitor expensive statements](#) link, the [Plan Trace](#) link, or the SAP HANA database explorer. You can also open the tool from the SAP Web IDE for SAP HANA. For more information on how to open the SQL Analyzer, see section *How can you access the SQL analyzer?* in *Analyzing Statement Performance*.

Procedure

1. Open the SQL Analyzer using the SAP HANA cockpit or the SAP Web IDE for SAP HANA.

The [SQL Analyzer](#) page opens, displaying the following views:

- Overview
 - Plan Graph
 - SQL
 - Operators
 - Statement Statistics (SQLScript only)
 - Timeline
 - Tables in Use
 - Table Accesses
2. Open the [Overview](#) tab to view important KPIs required to begin a performance analysis before going into the complex details.

KPI Description

KPI	Description
Time	The initial compilation time
	The total duration of the query excluding the compilation time
	The elapsed time indicating the total response time from the query execution request time to the end time of the query execution
Dominant Operators	Operators sorted by their execution time
Tables in Use	Total number of tables touched by any operators during execution
Result Record Count	The final result record count
Distribution	Number of SAP HANA index servers that are related to the query execution
Memory Allocated	Total memory allocated for executing the statement
System Version	The version of the system where the execution occurred

3. Open the [Plan Graph](#) tab to understand and analyze the execution plan of an SQL statement. It displays a visualization of a critical path based on inclusive execution time of operators and allows you to identify the most expensive path in a query execution plan.

In case of a SQLScript, the [Plan Graph](#) displays its complete definition. To retrieve the information in a text format, you can copy the definition by clicking the copy icon.

In the [Plan Graph](#) tab you can open the [Detail Properties](#) view by clicking one of the operators. This view offers detailed information on the operator such as name, location, ID, summary.

Furthermore, you can open the edge information detail by clicking one of the links between the operators. This view offers further information on the edge values, such as target, source, output cardinality, fetch call count, and output cardinality (estimation).

4. Open the [SQL](#) tab to get the complete view of the SQL statement string that is being analyzed.
5. Open the [Operators](#) tab to pinpoint specific operators of interest.

The view lists characteristics of all operators and supports:

- Display of various KPIs, for example physical (whether an operation is a real, physically executed one), offset, execution time, CPU time
- Setting of filters along all the columns/KPIs
- Display of the number of operators within the filtered set
- Immediate aggregated information (max, min, sum, and so on) for the same KPIs on the filtered operator set
- Detailed display of all operators within the filtered set

6. Open the [Tables in Use](#) tab for an overview of which tables have been used during the processing of a statement.

The view displays the following information:

- The name of the table
 - The host and port of the table's partition in the *Location* column
 - The partition number of the table's partition in the *Partition* column
 - The maximum possible amount of data processed for a given table during the execution of a statement, including the possibility of multiple accesses, in the *Max. Entries Processed* column
 - How often a table has been accessed during statement execution in the *Number of accesses* column
 - The maximum processing time across the possibly multiple table accesses in the *Maximum processing time* column
7. Open the *Statement Statistics* to view a set of statistics (SQLScript only) for each SQL statement involved in the procedure. This set of statistics provides a good starting point for analyzing the performance of a procedure as it lets users easily drill-down the most expensive SQL statement.

The following information is available for each statement so that users can sort the column (criterion) of their interest to find the most expensive statement: SQL Statement, Line Number, Execution Count, Execution Times, Compilation Times, Memory Allocated, Result Record Count, Explain Plan Result, and Procedure Comment.

8. Open the *Timeline* tab to get a complete overview of the execution plan based on the visualization of sequential time-stamps. The operator tree table displays hierarchical parent-child relationships and container-inner plan relationships. Based on the operators, the timeline chart shows the operations executed at different time intervals.
9. Open the *Table Accesses* tab to see the details on the table accesses performed during the processing of a statement.

The view displays the following information:

- The *Offset* time for accessing the table
- The name of the table
- The *Conditions* that affect the table accesses
- The *Processing Time*
- The amount of entries processed during an operation in the *Entries Processed* column
- The host and port of the table's partition in the *Location* column
- The partition number of the table's partition in the *Partition* column

Optional: To get aggregated information for each column, choose the aggregator functions in the drop-down menu under the column name. If an aggregator is chosen, you can see *More* information for your chosen query.

Optional: To refine results, click on *Filters* on the upper right corner. The *Filters* page appears, where you can choose the information type, operator, and enter values according to which you wish to filter the *Table Accesses*. Choose *OK* to confirm. If you want to remove the filters you selected, go to the *Filters* page and choose to *Restore* to defaults.

Optional: To sort the results, click the sorting icon next to the *Filters* button and choose the sorting order.

Optional: To customize the columns displayed in the information list, click on the settings icon on the upper right corner and choose to hide or display the desired columns.

10. Optional: Re-execute the SQL query by clicking the *Re-execute* button on the top right corner. If the query is parameterized, you can change the parameter values.

With a parameterized query, the *Input Parameters* appears, prompting you to enter your desired parameters in the *Parameter* tab. Where applicable, you can check the *Empty value* box. The values you

entered are reflected in the SQL string that you can see in the [SQL Statement](#) tab. Click on [Execute SQL](#) to analyze the SQL string according to the parameter values you entered and see the result on the [SQL Analyzer](#) page.

Related Information

[Analyzing Statement Performance \[page 440\]](#)

6.7.1.3.2.2 Monitor and Analyze Statements with Plan Trace

Plan Trace is a trace feature for SQL analyzer.

Context

Plan Trace enables you to collect SQL queries and their execution plans, executed in a given time frame. For each SQL query traced you can visualize the execution plan for performance analysis. Only 'SELECT' statements are traced with Plan Trace.

Procedure

1. On the [Overview](#) page, open the [Plan Trace](#) link from the [Administration](#) link list.
The [Plan Trace](#) page opens, displaying a set of statistics for each SQL statement collected in a given time frame. To find the most expensive SQL statements use the displayed categories (for example, start and end time, schema, user, statement hash).
2. Click the [Configure Trace](#) button on the bottom right to configure the plan trace.
The [Configuration](#) dialog opens, allowing you to set the options you want.
3. Optional: Open the selected statement string with the SQL analyzer by clicking the [More](#) link in the [Full SQL Statement](#) dialog.

6.7.1.3.2.3 Manage Saved Plans

The SQL analyzer result page shows SQL plans saved from a previously executed query.

Context

The saved plans feature allows you to revisit SQL statement queries that were executed with the SQL Analyzer in a previous session without having to re-execute them in the current session.

Procedure

1. On the *Overview* page, open the *Manage saved plans* link from the *DB Administration* link list.
The *SQL Analyzer result* page opens, displaying the *Saved Plans* table, containing the collected information on previously executed SQL queries. To find the your desired SQL statements, use the displayed categories (for example, system version, statement string, plan type, user name, schema name, statement hash, and so on).
2. Optional: You can delete saved plans by marking the checkbox on the left of the listed statement, and selecting *Delete*.
3. Optional: You can search saved plans by entering a statement string keyword in the search field.
4. Optional: You can sort the order of the table and filter results by the desired parameter.
5. Optional: You can customize what columns are shown in the table in the settings menu.

Related Information

[Analyzing Statement Performance \[page 440\]](#)

[Analyze Statement Performance \[page 442\]](#)

6.7.1.4 Improving Performance in SAP HANA Cockpit

You can improve the performance of the database using the SAP HANA cockpit.

You can use the following tools to analyze fine-grained aspects of system performance in the SAP HANA cockpit:

- Use the *Manage plan stability* to restore performance speed from the previous to the current system.
- Use the *Manage statement hints* to add statement hints to an SQL statement without modifying the actual statement in the application.

Related Information

[Managing Plan Stability \[page 447\]](#)

[Managing Statement Hints \[page 448\]](#)

6.7.1.4.1 Managing Plan Stability

Plan stability helps ensure the fast performance of queries by capturing query plans in a source system and reusing them in a target system to regenerate the original query plan.

In SAP HANA, the SQL query processor parses SQL statements and generates SQL query execution plans. As the query processor and the query optimizer continue to be developed (in, for example, the new HANA Execution Engine - HEX) the resultant execution plans for a given query may change from one HANA release to another. Although all developments are intended to improve performance of a query, it is possible it might not be equivalent after an upgrade.

In order to guarantee the performance of a query in new system upgrades, the plan stability feature offers the option to preserve a query's execution plan by capturing an abstraction of the plan and reusing it after the upgrade to regenerate the original plan and retain the original performance. In some cases, using statement hints may provide a solution to a loss of performance, see *Managing Statement Hints*.

Related Information

[Manage Plan Stability \[page 447\]](#)

[Managing Statement Hints \[page 448\]](#)

6.7.1.4.1.1 Manage Plan Stability

Use plan stability to capture query plans in a source system and reuse them in a target system to regenerate the original query plan.

Procedure

1. In the *DB Administration* list, choose the *Manage plan stability* link.

The *Manage Plan Stability* page opens. If you have any captured abstract SQL plans from the source system, they are displayed in the table view.

2. To capture SQL query plans from the source system, turn the *Capture Abstract SQL Plans* on.

You are prompted to configure the capture options. Filter the captures by user, choose whether you want to include plans from the SQL Plan Cache, and select *Start Capture*.

The [Capture Status](#) appears on the top of the screen, displaying the progress of captured execution plans and plans from SQL Plan Cache.

Optional: Press the refresh button next to the [Capture Status](#) to see the captured plans in the table below.

Optional: If you want to terminate the capture process prematurely, turn the [Capture Abstract SQL Plans](#) off. After the process is turned off, the captured query plans are displayed in the [Captured Plan](#) table.

3. Optional: See the full SQL statement string of a query by selecting [More](#) next to the visible statement string snippet in the table.
4. To apply the imported SQL query plans to the current system, turn the [Apply Abstract SQL Plans](#) on.

i Note

This will only apply to the SQL queries listed in the [Captured Plan](#) table.

Optional: Enable or disable the SQL query plans being applied to the current system by selecting them in the table and choosing [Enable](#) or [Disable](#) above the table. You can also [Select All](#) query plans, or choose to [Delete](#) your selections.

Optional: To see the progress of the query plan application, press the refresh button.

5. Optional: You can adjust the settings such as the maximum number of saved plans and maximum memory allocation in the settings dialog.

6.7.1.4.2 Managing Statement Hints

Use the [Manage statement hints](#) to add statement hints to an SQL statement without modifying the actual statement in the application.

The [Manage statement hints](#) allows you to pair an SQL statement string with a string of hints to be used during execution. Whenever a particular SQL statement is then executed in SAP HANA, the assigned statement hints are automatically added to the statement for execution.

You can find the [Manage statement hints](#) link in the [DB Administration](#) link list on the SAP HANA cockpit.

6.7.2 Monitoring and Analyzing Performance in SAP HANA Studio

Gathering and analyzing data regarding the performance of your SAP HANA systems is important for root-cause analysis and the prevention of future performance issues. The SAP HANA studio provides a number of tools for this purpose.

General information about overall system performance is available in the System Monitor and on the [Overview](#) tab of the Administration editor. You can monitor the following fine-grained aspects of system performance on the [Performance](#) tab:

- Threads
- Sessions
- Blocked transactions

- Execution statistics of frequently-executed queries in the SQL plan cache
- Expensive statements
- Progress of long-running jobs
- System load

Related Information

[Filters on the Performance Tab \[page 449\]](#)

[Thread Monitoring \[page 450\]](#)

[Session Monitoring \[page 452\]](#)

[Blocked Transaction Monitoring \[page 454\]](#)

[Monitoring SQL Performance with the SQL Plan Cache \[page 455\]](#)

[Expensive Statements Monitoring \[page 456\]](#)

[Job Progress Monitoring \[page 456\]](#)

[Load Monitoring \[page 457\]](#)

6.7.2.1 Filters on the Performance Tab

As the information available on the *Performance* tab is very detailed, some useful filters are available with which you can customize the amount and type of information displayed.

Filter	Description	Sub-Tab
Table Viewer	Use the table viewer to show and hide columns. To open the table viewer, choose  (<i>Configure Viewer</i>) in the toolbar on the top-right of the screen.	All
Hide Sessions	Use this filter to hide idle sessions, as well as sessions originating in the Administration editor or the SAP HANA studio.	<ul style="list-style-type: none"> • <i>Threads</i> • <i>Sessions</i> • <i>Blocked Transactions</i>
Column Filter	Use this filter to see information by distinct values in visible columns. To open the column filter, choose  <i>Filters...</i> in the toolbar on the top-right of the screen.	<ul style="list-style-type: none"> • <i>Sessions</i> • <i>SQL Plan Cache</i> • <i>Expensive Statements</i> • <i>Job Progress</i>
Host/Service/Thread Type	Use these filters to show threads from a specific host or service, or of a specific type.	<i>Threads</i>
Free-text filter	Use this filter to find a specific character string or expression in the displayed information.	<ul style="list-style-type: none"> • <i>Sessions</i> • <i>Blocked Transactions</i> • <i>SQL Plan Cache</i> • <i>Expensive Statements</i> • <i>Job Progress</i>

Any filters or layout configuration that you apply on the following sub-tabs are saved when you close the SAP HANA studio and applied the next time you open the tab. This is independent of which system you open.

- [Sessions](#)
- [Expensive Statements Trace](#)
- [SQL Plan Cache](#)
- [Job Progress](#)

6.7.2.2 Thread Monitoring

You can monitor all running threads in your system in the Administration editor on the [Performance > Threads >](#) sub-tab. It may be useful to see, for example, how long a thread is running, or if a thread is blocked for an inexplicable length of time.

Thread Display

By default, the [Threads](#) sub-tab shows you a list of all currently active threads with the [Group and sort](#) filter applied. This arranges the information as follows:

- Threads with the same connection ID are grouped.
- Within each group, the call hierarchy is depicted (first the caller, then the callee).
- Groups are displayed in order of descending duration.

On big systems with a large number of threads, this arrangement provides you with a more meaningful and clear structure for analysis. To revert to an unstructured view, deselect the [Group and sort](#) checkbox or change the layout in some other way (for example, sort by a column).

Thread Information

Detailed information available on the [Threads](#) sub-tab includes the following:

- The context in which a thread is used
This is indicated by the thread type. Important thread types are `SqlExecutor` and `PlanExecutor`. `SqlExecutor` threads handle session requests such as statement compilation, statement execution, or result fetching issued by applications on top of SAP HANA. `PlanExecutor` threads are used to process column-store statements and have an `SqlExecutor` thread as their parent.

i Note

With revision 56, `PlanExecutor` threads were replaced by `JobWorker` threads.

i Note

The information in the [Thread Type](#) column is only useful to SAP Support for detailed analysis.

- What a thread is currently working on
The information in *Thread Detail*, *Thread Method*, and *Thread Status* columns is helpful for analyzing what a thread is currently working on. In the case of *SqlExecutor* threads, for example, the SQL statement currently being processed is displayed. In the case of *PlanExecutor* threads (or *JobWorker* threads as of revision 56), details about the execution plan currently being processed are displayed.

i Note

The information in the *Thread Detail*, *Thread Method*, and *Thread Status* columns is only useful to SAP Support for detailed analysis.

- Information about transactionally blocked threads

A transactionally blocked thread is indicated by a warning icon (⚠) in the *Status* column. You can see detailed information about the blocking situation by hovering the cursor over this icon.

A transactionally blocked thread cannot be processed because it needs to acquire a transactional lock that is currently held by another transaction. Transactional locks may be held on records or tables. Transactions can also be blocked waiting for other resources such as network or disk (database or metadata locks).

The type of lock held by the blocking thread (record, table, or metadata) is indicated in the *Transactional Lock Type* column.

The lock mode determines the level of access other transactions have to the locked record, table, or database. The lock mode is indicated in the *Transactional Lock Type* column.

Exclusive row-level locks prevent concurrent write operations on the same record. They are acquired implicitly by update and delete operations or explicitly with the SELECT FOR UPDATE statement.

Table-level locks prevent operations on the content of a table from interfering with changes to the table definition (such as drop table, alter table). DML operations on the table content require an **intentional exclusive** lock, while changes to the table definition (DDL operations) require an exclusive table lock. There is also a LOCK TABLE statement for explicitly locking a table. Intentional exclusive locks can be acquired if no other transaction holds an exclusive lock for the same object. Exclusive locks require that no other transaction holds a lock for the same object (neither intentional exclusive nor exclusive).

For more detailed analysis of blocked threads, information about low-level locks is available in the columns *Lock Wait Name*, *Lock Wait Component* and *Thread ID of Low-Level Lock Owner*. Low-level locks are locks acquired at the thread level. They manage code-level access to a range of resources (for example, internal data structures, network, disk). Lock wait components group low-level locks by engine component or resource.

The *Blocked Transactions* sub-tab provides you with a filtered view of transactionally blocked threads.

Monitoring and Analysis Features

To support monitoring and analysis, you can perform the following actions on the *Threads* sub-tab:

- See the full details of a thread by right-clicking the thread and choosing *Show Details*.
- End the operations associated with a thread by right-clicking the thread and choosing *Cancel Operations*.

i Note

This option is not available for threads of external transactions, that is those with a connection ID of -1.

- Jump to the following related objects by right-clicking the thread and choosing **► Navigate To ► <related object> ►**:

- Threads called by and calling the selected thread
- Sessions with the same connection ID as the selected thread
- Blocked transactions with the same connection ID as the selected thread
- View the call stack for a specific thread by selecting the *Create call stacks* checkbox, refreshing the page, and then selecting the thread in question.

i Note

The information contained in call stacks is only useful to SAP Support for detailed analysis.

- Activate the expensive statements trace, SQL trace, or performance trace by choosing ► *Configure Trace* ► *<required trace>* ►.
- The *Trace Configuration* dialog opens with information from the selected thread automatically entered (application and user).

i Note

If the SQL trace or expensive statements trace is already running, the new settings overwrite the existing ones. If the performance trace is already running, you must stop it before you can start a new one.

Related Information

[SQL Trace \[page 672\]](#)

[Performance Trace \[page 675\]](#)

[Expensive Statements Trace \[page 676\]](#)

6.7.2.3 Session Monitoring

You can monitor all sessions in your landscape in the Administration editor on the ► *Performance* ► *Sessions* ► sub-tab.

Session Information

The *Sessions* sub-tab allows you to monitor all sessions in the current landscape. You can see the following information:

- Active/inactive sessions and their relation to applications
- Whether a session is blocked and if so which session is blocking
- The number of transactions that are blocked by a blocking session
- Statistics like average query runtime and the number of DML and DDL statements in a session
- The operator currently being processed by an active session (*Current Operator* column).

i Note

In earlier revisions, you can get this information from the SYS.M_CONNECTIONS monitoring view with the following statement:

```
SELECT CURRENT_OPERATOR_NAME FROM M_CONNECTIONS WHERE CONNECTION_STATUS =  
'RUNNING'
```

→ Tip

To investigate sessions with the connection status RUNNING, you can analyze the SQL statements being processed in the session. To see the statements, ensure that the *Last Executed Statement* and *Current Statement* columns are visible. You can then copy the statement into the SQL console and analyze it using the *Explain Plan* and *Visualize Plan* features. It is also possible to use the SQL plan cache to understand and analyze SQL processing.

Monitoring and Analysis Features

To support monitoring and analysis, you can perform the following actions on the *Sessions* sub-tab:

- Cancel a session by right-clicking the session and choosing *Cancel Session...*
- Jump to the following related objects by right-clicking the session and choosing **▶ Navigate To > <related object> ▾**:
 - Threads with the same connection ID as the selected session
 - Blocked transactions with the same connection ID as the selected session
- Activate the performance trace, SQL trace, or expensive statements trace by choosing **▶ Configure Trace > <required trace> ▾**.

The *Trace Configuration* dialog opens with information from the selected session automatically entered (application and user).

i Note

If the SQL trace or expensive statements trace is already running, the new settings overwrite the existing ones. If the performance trace is already running, you must stop it before you can start a new one.

Related Information

[SQL Trace \[page 672\]](#)

[Performance Trace \[page 675\]](#)

[Expensive Statements Trace \[page 676\]](#)

6.7.2.4 Blocked Transaction Monitoring

Blocked transactions, or transactionally blocked threads, can impact application responsiveness. They are indicated in the Administration editor on the **Performance > Threads** tab. You can see another representation of the information about blocked and blocking transactions on the *Blocked Transactions* sub-tab.

Information About Blocked Transactions

Blocked transactions are transactions that are unable to be processed further because they need to acquire transactional locks (record or table locks) that are currently held by another transaction. Transactions can also be blocked waiting for other resources such as network or disk (database or metadata locks).

The type of lock held by the blocking transaction (record, table, or metadata) is indicated in the *Transactional Lock Type* column.

The lock mode determines the level of access other transactions have to the locked record, table, or database. The lock mode is indicated in the *Transactional Lock Type* column.

Exclusive row-level locks prevent concurrent write operations on the same record. They are acquired implicitly by update and delete operations or explicitly with the SELECT FOR UPDATE statement.

Table-level locks prevent operations on the content of a table from interfering with changes to the table definition (such as drop table, alter table). DML operations on the table content require an **intentional exclusive** lock, while changes to the table definition (DDL operations) require an exclusive table lock. There is also a LOCK TABLE statement for explicitly locking a table. Intentional exclusive locks can be acquired if no other transaction holds an exclusive lock for the same object. Exclusive locks require that no other transaction holds a lock for the same object (neither intentional exclusive nor exclusive).

For more detailed analysis of blocked transactions, information about low-level locks is available in the columns *Lock Wait Name*, *Lock Wait Component* and *Thread ID of Low-Level Lock Owner*. Low-level locks are locks acquired at the thread level. They manage code-level access to a range of resources (for example, internal data structures, network, disk). Lock wait components group low-level locks by engine component or resource.

Monitoring and Analysis Features

To support monitoring and analysis, you can perform the following actions on the *Blocked Transactions* sub-tab:

- Jump to threads and sessions with the same connection ID as a blocked/blocking transaction by right-clicking the transaction and choosing **Navigate To > <related object>**.
- Activate the performance trace, SQL trace, or expensive statements trace for the blocking transaction (that is the lock holder) by choosing **Configure Trace > <required trace>**. The *Trace Configuration* dialog opens with information from the selected thread automatically entered (application and user).

i Note

If the SQL trace or expensive statements trace is already running, the new settings overwrite the existing ones. If the performance trace is already running, you must stop it before you can start a new one.

Related Information

[SQL Trace \[page 672\]](#)

[Performance Trace \[page 675\]](#)

[Expensive Statements Trace \[page 676\]](#)

6.7.2.5 Monitoring SQL Performance with the SQL Plan Cache

The SQL plan cache can provide you with an insight into the workload in the system as it lists frequently executed queries. You can view the plan cache in the Administration editor on the **Performance** > **SQL Plan Cache** sub-tab.

Technically, the plan cache stores compiled execution plans of SQL statements for reuse, which gives a performance advantage over recompilation at each invocation. For monitoring reasons, the plan cache keeps statistics about each plan, for instance number of executions, min/max/total/average runtime, and lock/wait statistics. Analyzing the plan cache is very helpful as one of the first steps in performance analysis because it gives an overview about what statements are executed in the system.

i Note

Due to the nature of a cache, seldom-used entries are evicted from the plan cache.

The SQL plan cache is useful for observing overall SQL performance as it provides statistics on compiled queries. Here, you can get insight into frequently executed queries and slow queries with a view to finding potential candidates for optimization.

The following information may be useful:

- Dominant statements (TOTAL_EXECUTION_TIME)
- Long-running statements (AVG_EXECUTION_TIME)
- Frequently-executed plans (EXECUTION_COUNT)
- Number of records returned (TOTAL_RESULT_RECORD_COUNT)
- Statements with high lock contention (TOTAL_LOCK_WAIT_COUNT)

i Note

The collection of SQL plan cache statistics is enabled by default, but you can disable it on the **SQL Plan Cache** tab by choosing *Configure*.

To help you understand and analyze the execution plan of an SQL statement further, you can generate a graphical view of its plan by right-clicking the statement and choosing [Visualize Plan](#).

The system views associated with the SQL plan cache are M_SQL_PLAN_CACHE_OVERVIEW and M_SQL_PLAN_CACHE.

6.7.2.6 Expensive Statements Monitoring

Expensive statements are individual SQL queries whose execution time was above a configured threshold. The expensive statements trace records information about these statements for further analysis and displays them in the Administration editor on the [Performance > Expensive Statements Trace](#) sub-tab.

The individual steps of statement execution are displayed in a hierarchical tree structure underneath aggregated statement execution information.

The following information may be useful:

- When the query started (START_TIME)
- How long the query took (DURATION_MICROSEC)
- Name(s) of the objects accessed (OBJECT_NAME)
- The SQL statement (STATEMENT_STRING)

i Note

The expensive statements trace is deactivated by default. You can activate and configure it either here on the [Expensive Statements Trace](#) sub-tab, or on the [Trace Configuration](#) tab.

To help you understand and analyze the execution plan of an expensive statement further, you can generate a graphical view of its plan by right-clicking the statement and choosing [Visualize Plan](#).

Related Information

[Expensive Statements Trace \[page 676\]](#)

6.7.2.7 Job Progress Monitoring

Certain operations in SAP HANA typically run for a long time and may consume a considerable amount of resources. You can monitor long-running jobs in the Administration editor on the [Performance > Job Progress](#) sub-tab.

By monitoring the progress of long-running operations, for example, delta merge operations and data compression, you can determine whether or not they are responsible for current high load, see how far along they are, and when they will finish.

The following information is available, for example:

- Connection that triggered the operation (CONNECTION_ID)
- Start time of the operation (START_TIME)
- Steps of the operation that have already finished (CURRENT_PROGRESS)
- Maximum number of steps in the operation (MAX_PROGRESS)

For more information about the operations that appear on the *Job Progress* sub-tab, see system view M_JOB_PROGRESS.

6.7.2.8 Load Monitoring

A graphical display of a range of system performance indicators is available in the Administration editor on the **► Performance ► Load ►** sub-tab.

You can use the load graph for performance monitoring and analysis. For example, you can use it to get a general idea about how many blocked transactions exist now and in the past, or troubleshoot the root cause of slow statement performance.

6.7.3 Performance: Using Hints to Query Data Snapshots

Several features in SAP HANA use data snapshots to improve performance. You can use configurable hint classes as a standard way of controlling at run time how the data is selected, either from the snapshot or from the database.

Several features in SAP HANA use data snapshots or replicated data to improve performance; this includes:

- Result Cache
- Asynchronous Table Replication
- System Replication (Active/Active Read Enabled).

Snapshots carry the risk of holding stale data and administrators need to be able to specify a maximum time value up to which the replicated data may lag behind the live data. The features listed above all use the RESULT_LAG() hint with a set of standard configurable hint classes as a common way of controlling how the data is selected at run time (also referred to as hint-based routing). Hint classes give the administrator a tool to balance a system in terms of query response time, query load, resource utilization and freshness of data. Moreover, they de-couple SAP HANA features from application development and administrator choices.

Using the Hint RESULT_LAG()

The RESULT_LAG() hint takes two parameter values to determine:

- Which is the preferred data source (snapshot or live data)
- How much time the retrieved data may lag behind the live data.

The syntax for the RESULT_LAG() hint is as shown here:

```
WITH HINT (RESULT_LAG("<short_class_name>", [seconds]))
```

Class name: the following pre-defined hint classes are available for use with this hint, these, and options for configuring them, are described in detail below. Only the last two elements of the name are used in SQL statements, for example 'hana_cache', 'hana_short' and so on:

- hint_result_lag_hana_cache - for Result Cache, the default lag time value is the cache retention time
- hint_result_lag_hana_atr - for Asynchronous Table Replication, the default lag time value can be defined as a configuration parameter
- hint_result_lag_hana_sr - for System Replication (Active/Active), the default lag time value can be defined as a configuration parameter
- hint_result_lag_hana_short
- hint_result_lag_hana_long

Note that customers and developers may also add classes of their own (and corresponding configuration values) but these must not use the hana_* or sap_* naming convention.

Seconds: This parameter is an optional time value in seconds; this is the maximum time lag up to which the data is judged to be acceptable; if the snapshot or replica is older then it will not be used and the query is routed to the data source. If no seconds value is given the default value will be used. In some cases this feature can be disabled by setting a configuration option - see details below.

Configuration of Routing

Each hint class has its own section in the indexserver.ini file ([hana_result_lag_hana_short] for example, sections must also be added for user-defined classes) where one or more parameters are available. All classes have the `enable_features` parameter; this is the 'routing' parameter which is used to define an order of preference for the data source used. The following table shows the default `enable_features` setting for each class, for example, the *hana_long* class is defined by default with the sequence, 'resultcache,atr,sr', meaning: data from the resultcache is preferred, if not then data from a replica table (asynchronous table replication), otherwise data from a secondary system (system replication). The same maximum lag time value defined is applied to the whole sequence of data sources (an example is given below):

Class	Parameter	Default Setting (Preferred Source Sequence)
hana_cache	enable_features	resultcache
hana_atr	enable_features	atr
hana_sr	enable_features	sr
hana_short	enable_features	atr, sr
hana_long	enable_features	resultcache,atr,sr

More detail of how these classes are used in each of the three application areas is described in the following subsections.

Hint Class for Result Cache

Using the class `hana_cache` the result cache of the target view / table function is exploited if it is enabled and applicable. This is shown in the following example:

```
SELECT * FROM v1 WITH HINT( RESULT_LAG ( 'hana_cache', 30 ) );
```

If a value for the seconds (time) parameter of the hint is given it overrides any retention time value which has been defined for the view. For example, the 30 second time lag defined in the above example would override the hundred minute retention period defined when the cache was created:

```
ALTER VIEW v1 ADD CACHE RETENTION 100;
```

The result cache and dynamic result cache are described in detail in the *SAP HANA Troubleshooting and Performance Analysis Guide*.

Hint Class for Asynchronous Table Replication

You can use the `hana_atr` class to access replicated tables.

```
SELECT * FROM T1 WITH HINT( RESULT_LAG ('hana_atr') , 10 );
```

If the current lag time for the data on the replica is less than the stated value for the [seconds] parameter (10 seconds in this example) then the query is executed on the replicated table, otherwise it would be executed on the source.

More Configuration Options for Class `hana_atr`

Parameter:	<code>atr_default_lag_time</code>
Purpose	The default lag time value for this class. This value will be applied if no seconds value is entered with the hint on the SQL command line.
Default	-1 In this case if no seconds value is entered with the hint the query is always routed to the replica.
Unit	seconds

Parameter:	<code>atr_ignore_lag_time</code>
Purpose	Set this to true to disable the check for the specified lag time so that the query will always access replica tables regardless of the value entered with the hint and the default lag time setting.
Default	false

For more information refer to the *Asynchronous Table Replication* section of this guide.

Hint Class for System Replication (Active/Active read enabled)

You can use the class `hana_sr` for system replication and select data from the secondary system. This is shown in the following example:

```
SELECT * FROM T1 WITH HINT( RESULT_LAG ('hana_sr', 60) );
```

If this is used in a situation where the result cache might also be available, the following table shows specifically which data source is used in either of the two possible cases (result cache available yes or no):

enable_features	Result Cache Available?	Data Source
sr,resultcache	Yes	Secondary (cache used)
sr,resultcache	No	Secondary
resultcache,sr	Yes	Primary (cache used)
resultcache,sr	No	Secondary
sr	Yes	Secondary
sr	No	Secondary
resultcache	Yes	Primary (cache used)
resultcache	No	Primary

This table shows the behavior on the primary system. If the application connects to the secondary system directly the result cache can be also used at secondary. For more information refer to the *Active/Active (Read Enabled)* section of this guide.

More Configuration Options for Class `hana_sr`

Other parameters are also available for the `hana_sr` hint class:

Parameter:	<code>sr_default_lag_time</code>
Purpose	The default lag time value for this class. This value will be applied if no value is entered with the hint on the SQL command line.
Default	-1 In this case if no seconds value is entered with the hint the query is always routed to the primary.
Unit	seconds
Parameter:	<code>sr_ignore_lag_time</code>
Purpose	Set this to true to disable the check for the specified lag time.
Default	false

Parameter:	sr_enable_primary_redirection_for all errors
Purpose	The routed statement will be redirected to primary for all errors.
Default	true
Parameter:	sr_enable_primary_redirection
Purpose	This parameter is used for error handling: if an out of memory error occurs on the secondary, or if the lag time is exceeded, then the routed statement will be redirected to the primary and executed.
Default	true

Examples Using HANA_SHORT and HANA_LONG

You can use `hana_short` and `hana_long` classes in combination with other application-specific hint classes. Here, we give two examples; in the first a value for the `seconds` parameter is given, and in the second the default values for each data source as specified in the configuration file are used. The configuration file for `hana_long` for example, could be as follows:

Sample Code

```
[hana_result_lag_hana_long]
enable_features = resultcache,atr,sr
atr_default_lag_time = 10
sr_default_lag_time = 20
```

The `enable_features` parameter specifies the sequence of data sources as: result cache preferred, secondly table replica, otherwise system replication secondary system.

Example 1 In this example a `seconds` value is given in the query and this value is used in all three cases to evaluate whether the cached/replicated data is acceptable.

```
SELECT * FROM V1 WITH HINT( RESULT_LAG ('hana_long') , 10 );
```

First if the result cache data is less than the `seconds` parameter value (10 seconds) the query will read data from here. If the result cache is too stale then the age of the replicated table data will be evaluated against the `seconds` parameter value and if this data is too stale then the system replicated data will be evaluated to determine whether to execute the query on the primary or secondary system. If none of the cached sources is available (`seconds` value smaller than the current lag time) then the primary system is accessed.

Example 2 In the following example (referring to the same sequence of three data sources) a time value is not given in the query, and the default lag values for each data source are used to evaluate whether the data source is acceptable:

```
SELECT * FROM V1 WITH HINT( RESULT_LAG ('hana_long') );
```

The default values are: for result cache the predefined retention period, for `atr` and `sr` the ini parameter values defined in the `[hana_result_lag_hana_long]` section apply, that is, `atr_default_lag_time` for asynchronous

replication (10 seconds in the sample code above) and `sr_default_lag_time` for system replication (20 seconds in the sample code above).

Related Information

[Active/Active \(Read Enabled\) \[page 1157\]](#)

[Asynchronous Table Replication Operations \[page 589\]](#)

6.7.4 Persistent Data Storage in the SAP HANA Database

Persistent storage media are required for ongoing save operations for data and redo log files.

To protect data from the risk of memory failure SAP HANA persists in-memory data to storage media and flushes all changed data from memory to the data volumes. This operation takes place on the basis of savepoints which occur by default every 5 minutes.

Customers have a choice of storage media types: in addition to conventional hard disk storage, non-volatile RAM is also supported for data storage, specifically, for the MAIN fragment of column store tables. For more information about the MAIN and DELTA fragments of tables refer to Delta Merge.

The following subsections describe how these types of persistence work and how they must be configured for use. The two storage types operate in a similar way but there are some essential differences and a separate section is dedicated to each type.

Related Information

[The Delta Merge Operation \[page 526\]](#)

6.7.4.1 Data and Log Volumes

To ensure that the database can always be restored to its most recent committed state, changes to data in the database are periodically copied to disk, logs containing data changes and certain transaction events are also saved regularly to disk. Data and logs of a system are stored in volumes.

SAP HANA persists in-memory data by using savepoints. Each SAP HANA service has its own separate savepoints. During a savepoint operation, the SAP HANA database flushes all changed data from memory to the data volumes. The data belonging to a savepoint represents a consistent state of the data on disk and remains so until the next savepoint operation has completed. Redo log entries are written to the log volumes for all changes to persistent data. In the event of a database restart (for example, after a crash), the data from the last completed savepoint can be read from the data volumes, and the redo log entries written to the log volumes since the last savepoint can be replayed.

The frequency at which savepoints are defined can be configured in the `persistence` section of the `global.ini` file (every 5 minutes by default). Savepoints are also triggered automatically by a number of other operations such as data backup, and database shutdown and restart. You can trigger a savepoint manually by executing the following statement `ALTER SYSTEM SAVEPOINT`.

You must always ensure that there is enough space on the disk to save data and logs. Otherwise, a disk-full event will occur and the database will stop working.

Directory Hierarchy for Data and Log Storage

During the installation process, the following default directories are created as the storage locations for data and log volumes:

- `/usr/sap/<SID>/SYS/global/hdb/data`
- `/usr/sap/<SID>/SYS/global/hdb/log`

i Note

These default directories are defined in the parameters `basepath_datavolumes` and `basepath_logvolumes` in the `persistence` section of the `global.ini` file.

These directories contain a separate sub-directory, or storage partition, for each host in the system. These are named `mnt00001`, `mnt00002`, `mnt00003` and so on, by default. Each host storage partition contains a sub-directory for every database service that persists data. These sub-directories represent the actual volumes. They are named `hdb00001`, `hdb00002`, `hdb00003`, and so on by default.

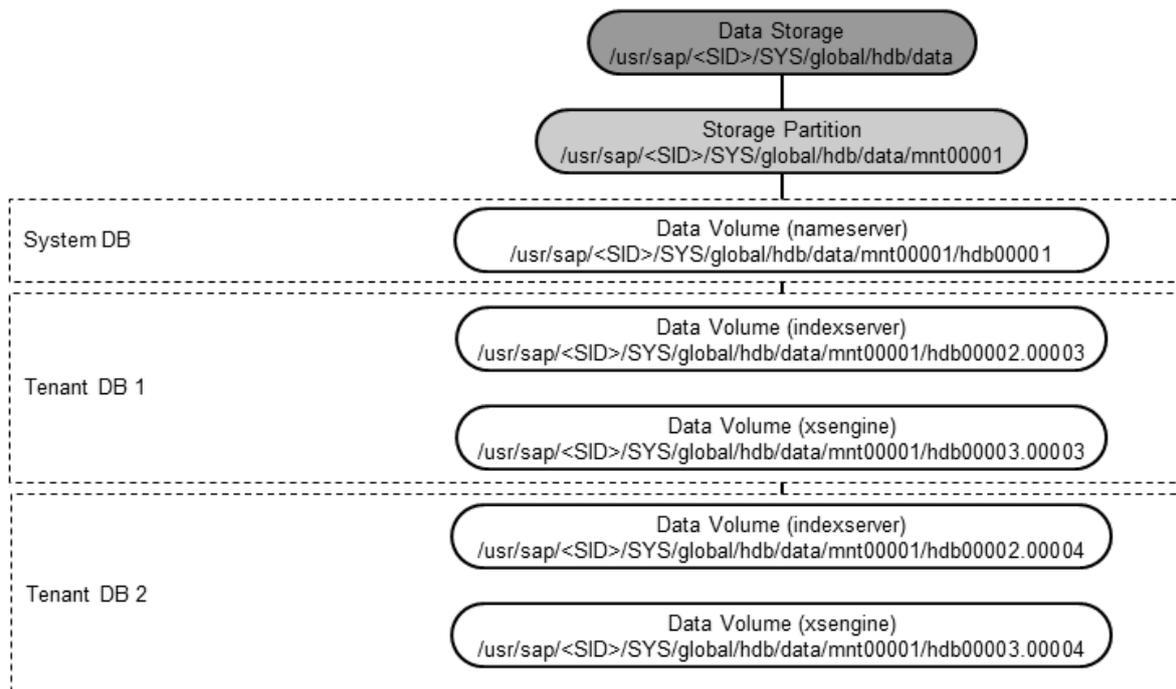
The volume names of tenant databases have a suffix to represent the database. For example, the indexserver volume for the first tenant database is `hdb00002.00002`, for the second database `hdb00002.00003`, and so on.

The services that persist data and therefore have volumes are the following:

Service	Note
nameserver	Only the nameserver service in the system database on the active master host persists data. Slave nameserver hosts communicate with the master, but do not persist data.
indexserver	The indexserver service in tenant databases on all hosts except standby hosts persists data. This data volume can be partitioned to improve performance when reading and writing to this server.
xsengine (if running as a separate service)	The xsengine service persists data on any host on which it is running.

The following figure illustrates the default storage hierarchy described above using the example of data storage. The system is a system with two tenant databases on a single host.

Directory Hierarchy for Persistent Data Storage



Data and Log Volumes

Data Volumes

By default each data volume contains one file (`datavolume_0000.dat`) in which data is organized into pages, ranging in size from 4KB to 16MB (page size class). Data is written to and loaded from the data volume page-wise. Over time, pages are created, changed, overwritten, and deleted. The size of the data file is automatically increased as more space is required. However, it is not automatically decreased when less space is required. This means that at any given time, the actual payload of a data volume (that is the cumulative size of the pages currently in use) may be less than its total size. This is not necessarily significant – it simply means that the amount of data in the file is currently less than at some point in the past (for example, after a large data load). The `Total Size` and the `Fill Ratio` values in the monitoring view `M_DATA_VOLUME_STATISTICS` give details about capacity.

If a data volume has a considerable amount of free space, it might be appropriate to shrink the data volume. However, a data file that is excessively large for its typical payload can also indicate a more serious problem with the database. SAP support can help you analyze your situation.

The indexserver data volume can be partitioned so that data can be segmented and distributed over several *stripes* on the disk. SQL commands are available to create and drop partitions, a numeric partition ID value is assigned automatically by the system. In this case, the volume ID and an offset value to identify the stripe where the data is saved are included in the volume file name (`datavolume_nnnn_<offset>.dat`). See *Partitioning Data Volumes* for details.

Log Volumes

Each log volume contains the file `logsegment_<partition_ID>_directory.dat` and one or more log segment files (`logsegment_<partition_ID>_<segment_number>.dat`). Currently only one log partition is

supported for each service, so the default file names are `logsegment_000_directory.dat` and `logsegment_000_00000000.dat`, `logsegment_000_00000001.dat`, `logsegment_000_00000002.dat` and so on. Log segment files are cyclically overwritten depending on the log mode. The log mode determines how logs are backed up. Log volumes only grow if there are no more segment files available for overwriting. Log segment files that are available for overwriting have been backed up and are not required for a database restart. If necessary you can remove these files to free up space in your file system by executing the SQL statement `ALTER SYSTEM RECLAIM LOG`. Note that new log segment files will need to be created later and this will affect performance.

You can monitor disk usage in SAP HANA cockpit using the Disk Usage and Performance Monitor apps. In SAP HANA studio you can monitor disk usage, volume size, and other disk activity statistics on the Volumes tab of the Administration editor. See *Monitoring Disk Space* and *Monitoring in SAP HANA Cockpit* for details.

⚠ Caution

Do not remove either data files or log files using operating system tools as this will corrupt the database.

Multiple Files in Data Volumes for Ext3 File Systems

The Ext3 file system has a file size limitation of 2TB. If the existing files in a data volume located in an Ext3 file system reach the 2 TB limit, SAP HANA automatically creates additional files. This allows the use of Ext3 file systems even with applications (in particular, single-host SAP ERP systems) that have a larger memory requirement per host.

For more information about splitting data backups into files of limited size, see the section on data backups.

Related Information

[Data Backups \[page 1246\]](#)

[Monitoring Disk Space \[page 378\]](#)

[Partitioning Data Volumes \[page 465\]](#)

[Monitoring in SAP HANA Cockpit \[page 317\]](#)

6.7.4.2 Partitioning Data Volumes

Data volumes on the indexserver can be partitioned so that read and write operations can run in parallel with increased data throughput.

Each SAP HANA service has one data volume which by default stores its data in one single file. Administrators have the option to partition the data volume for the indexserver so that data files can be segmented and distributed over several *stripes*. Partitioning is a logical distribution of data pages over different partitions and striping is a physical segmentation of files; using partitioning and striping data can be read or written in parallel threads.

If you create a new partition on the index server it is added simultaneously to all indexservers in the topology. New partitions become active after the next savepoint on each indexserver, this is shown in the partition STATE value which changes from ACTIVATING to ACTIVE. By default all data volumes have a single partition with the numeric ID zero. A numeric partition ID is assigned automatically to new partitions by the HANA persistency layer. If the partition numbering is for any reason inconsistent across all indexservers then any attempt to add new partitions will fail.

Partition details are saved in the indexserver.ini configuration file:

```
basepath_datavolume[<partition id>]=<path>
```

The `path` value is the default HANA data volume basepath; this cannot be changed but more flexibility will be introduced in future releases.

Details of the SQL statements to manage data volume partitions are given in the *SAP HANA SQL and System Views Reference*, the two main statements to add and drop partitions are shown here. If you execute the command to add a partition on the indexserver it will add a new partition to all indexservers in the topology:

```
ALTER SYSTEM ALTER DATAVOLUME ADD PARTITION
```

You can drop specific data volume partitions by specifying the partition ID number:

```
ALTER SYSTEM ALTER DATAVOLUME DROP PARTITION <id>
```

This command drops the identified partition from all indexservers in the topology. The default partition with ID zero cannot be dropped. If you drop a partition then all data stored in the partition is automatically moved to the remaining partitions and for this reason dropping a partition may take time. This operation also removes the partition entry from the configuration file.

You can see the current data volume configuration from the following two views:

- `M_DATA_VOLUME_STATISTICS`: This provides statistics for the data volume partitions on the indexserver including the number of partitions and size.
- `M_DATA_VOLUME_PARTITION_STATISTICS`: This view provides statistics for the individual partitions, identified by `PARTITION_ID`, and includes the partition STATE value.

6.7.4.3 Persistent Memory

Persistent memory (non-volatile RAM, also referred to as Storage Class Memory) is supported in SAP HANA as a persistent storage type.

Persistent memory (or NVRAM) is an emerging class of memory which combines the qualities of both DRAM and Flash storage and bridges the gap between disk storage and main memory. For the purposes of SAP HANA its most important characteristics are that it is byte addressable like DRAM and can be treated by the CPU as RAM therefore offering fast read and write performance; its latency characteristics are also very close to DRAM.

The latest information about this feature, such as hardware availability and operating system support, is provided in SAP Note 2618154 *SAP HANA Persistent Memory (NVM)*.

The persistent storage functions as DRAM memory and is used specifically for MAIN data fragments of the column store where approximately 90% of the data of a table is held. In this way DELTA data can continue to be

stored in DRAM (a relatively small quantity of data requiring the fastest possible access) with the MAIN segment of a table in NVRAM. Note that if persistent memory storage is enabled, data is still written to the data volumes but is simply ignored. After a system restart the main data fragments saved in persistent storage are still available in memory and do not have to be reloaded. This storage type has no impact on backup and recovery or high availability processes.

As with the data volume persistent storage feature, data is persisted on the basis of savepoints. In the case of non-volatile memory storage, NVM blocks are constructed whenever the in-memory MAIN fragment is created and are kept consistent according to HANA's savepoints.

The key performance benefit that this will bring is in greatly accelerated start-up times so that data can be quickly reloaded into memory and any system downtime can be minimized. Other benefits may include higher memory capacities per server.

Performance differences of this storage type compared to DRAM due to higher latency are mitigated by features such as hardware prefetchers. It is also possible to enable persistent memory storage for selected tables or for specific columns of a table so that a selective implementation of this storage type is possible (see Configuration below).

Installation

The persistent memory feature can be set up during system installation using the SAP HANA database lifecycle manager (HDBLCM) which uses two parameters to, firstly, enable persistent memory (`use_pmem`) and secondly the `pmempath` parameter to set the basepath (see Configuration below). For more information on lifecycle manager refer to the *SAP HANA Server Installation and Update Guide*.

At the level of the Linux operating system, SAP HANA only supports DAX enabled (Direct Access) file systems.

Filesystem layout of persistent memory block storage

The directory specified as the basepath (see Configuration) is the root persistent memory storage directory beneath which data for each savepoint is stored. The structure beneath the basepath is maintained by HANA and no attempt should be made to modify this content. Multiple logical versions of blocks are saved and are organized in subdirectories named on the basis of, for example, service name, Volume ID and current savepoint version. There are two subdirectory systems, one for data and one for deletion requests: the memory storage persistence maintains a history of requests to delete blocks in a *tombstones* directory parallel to the *data* directory. In the case of tombstone blocks the savepoint version relates to the savepoint when the data block was requested for deletion.

Configuration

The default behavior for using persistent memory is determined automatically: the system tests for the DAX filesystem at the defined persistent memory basepath. If the hardware is installed and the basepath correctly configured pointing to the persistent memory mount then all tables will use persistent memory by default. The SAP HANA installer sets the persistent memory basepath parameter if it is able to identify the DAX filesystem and if the user of the installation tool confirms the usage; the basepath parameter can also be set manually if necessary.

This default behavior can be overridden at four levels in the sequence shown here:

4	Database	Can be enabled or disabled by configuration parameter.
3	Table	Can be enabled or disabled by SQL statement.
2	Partition	Can be enabled or disabled by SQL statement.
1	Column	Can be enabled or disabled by SQL statement.

This level of granularity offers a high degree of flexibility: if persistent memory is applied at a certain level it is automatically inherited at lower levels of the sequence but can also be overridden at lower levels.

At the highest configuration level (database), persistent memory is managed by setting the `table_default` configuration parameter for the database as a whole. You can switch `table_default` to OFF to enable more selective persistent memory storage at lower levels for particular tables and partitions.

At other levels persistent memory is managed at the SQL command line using the `alter_persistent_memory_spec` clause with the CREATE / ALTER table commands. Here, we give examples to show how this can be applied at each level; refer to the *SAP HANA SQL and System Views Reference* for full details of this clause.

Configuration Parameter: Basepath

As with the other existing data and log persistent storage volumes, the storage location basepath for persistent memory storage must be defined as a configuration parameter in the `persistence` section of the `global.ini` file. Enter the basepath location in the following parameter:

Parameter	<code>basepath_persistent_memory_volumes</code>
Short Description	Location of NVRAM storage.
Full Description	Data for NVRAM-enabled tables or columns is stored at the location defined in this parameter. Multiple locations can be defined here using a semi-colon as a separator (no spaces). Multiple paths correspond to NUMA nodes without any ordering restriction in terms of NUMA node mapping, for example, for a four-NUMA node system: <code>"/mnt/pm0;/mnt/pm1;/mnt/pm4;/mnt/pm2"</code>
Type	Text
Change	Restart required (offline changeable only)
Default	Blank

Configuration Parameter: Enable persistent memory storage

The `table_default` parameter is in the `persistent_memory` section of the `indexserver.ini` file.

Parameter	<code>table_default</code>
Short Description	Enable or disable persistent memory storage for the database generally.

Parameter	table_default
Full Description	This parameter can be set to ON or OFF or DEFAULT (no explicit preference, follow the existing default behavior). The setting applies for the complete database but can be overridden by settings applied at a lower level.
Change	Online
Default	DEFAULT

SQL Alter Persistent Memory Clause

Persistent memory can be enabled by SQL command using the `alter_persistent_memory_spec` clause with CREATE and ALTER table. This can be used at any of the three levels - table, partition or column to enable or disable persistent memory storage using the PERSISTENT MEMORY switch. The following examples illustrate this using CREATE TABLE.

Example 1: Create table with persistent memory storage enabled for the new table.

```
CREATE COLUMN TABLE PMTABLE (C1 INT, C2 VARCHAR (10)) PERSISTENT MEMORY ON
```

PERSISTENT MEMORY can be set to ON, OFF or DEFAULT meaning no preference.

Example 2: Create range-partitioned table with persistent memory storage for a selected partition (this is only valid for range partitions):

```
CREATE COLUMN TABLE PMTABLE (C1 INT) PARTITION BY RANGE (C1) (
    PARTITION '0' <= VALUES < '10' PERSISTENT MEMORY ON,
    PARTITION OTHERS PERSISTENT MEMORY OFF);
```

Example 3: Create table with persistent memory storage applied to selected columns of a table:

```
CREATE COLUMN TABLE PMTABLE (C1 INT PERSISTENT MEMORY ON, C2 VARCHAR (10), C3
INT PERSISTENT MEMORY OFF)
```

For ALTER TABLE a preference value is also required (in addition to the ON/OFF switch) to determine how the change will be applied. This is the `alter_persistent_memory_preference` clause which requires one (or more) of the following keywords:

- IMMEDIATE - the change is applied immediately and the specified storage is populated.
- DEFERRED - the specified storage will be populated at the time of the next delta merge or reload.
- CASCADE - this keyword can be used at table level in addition to IMMEDIATE and DEFERRED to apply the change to all lower levels of the hierarchy.

The following example shows this usage with ALTER TABLE and cascades the change immediately to partitions and columns of the named table:

```
ALTER TABLE MYTABLE PERSISTENT MEMORY ON IMMEDIATE CASCADE
```

Monitoring Views

For full details of the monitoring views and values named here refer to the *SAP HANA SQL and System Views Reference*.

The following monitoring views show information about persistent memory in the `PERSISTENT_MEMORY` column: `M_CS_TABLES`, `M_CS_COLUMNS` and `M_CS_ALL_COLUMNS`. This shows either `TRUE` (persistent memory is enabled) or `FALSE`.

At the level of tables (`M_CS_TABLES`) usage information is shown as `PERSISTENT_MEMORY_SIZE_IN_TOTAL` and `PERSISTENT_MEMORY_SIZE_IN_MAIN`.

At the level of columns (`M_CS_COLUMNS` and `M_CS_ALL_COLUMNS`) usage information is shown in a series of columns named: `MAIN_PERSISTENT_MEMORY_SIZE_IN_*` These columns show usage in the categories `DATA`, `DICT`, `INDEX`, `MISC` and `UNUSED`. These column views also include `STORED_IN_PERSISTENT_MEMORY` and `LOADED_FROM_PERSISTENT_MEMORY`, both `TRUE/FALSE` values.

Three monitoring views show full statistical details of persistent memory usage:

- `M_PERSISTENT_MEMORY_VOLUMES` - capacity, usage and metadata of persistent memory volumes configured per NUMA Node.
- `M_PERSISTENT_MEMORY_VOLUME_DATA_FILES` - metadata statistics about files created by SAP HANA services for data storage on the persistent memory volumes.
- `M_PERSISTENT_MEMORY_VOLUME_STATISTICS` - statistics of physical lifecycle events of blocks managed by SAP HANA services on the persistent memory volumes.

Related Information

[SAP Note 2618154](#)

6.7.5 Memory Usage in the SAP HANA Database

Memory is a fundamental resource of the SAP HANA database. Understanding how the SAP HANA database requests, uses, and manages this resource is crucial to the understanding of SAP HANA.

SAP HANA provides a variety of memory usage indicators that allow for monitoring, tracking, and alerting. The most important indicators are used memory and peak used memory. Since SAP HANA contains its own memory manager and memory pool, external indicators such as the size of resident memory at host level and the size of virtual and resident memory at process level can be misleading when you are estimating the real memory requirements of an SAP HANA deployment.

For more information about memory consumption with regards to SAP HANA licenses, see SAP Note 1704499.

Related Information

[SAP HANA Used Memory \[page 269\]](#)

[Memory Sizing \[page 271\]](#)

[Allocated Memory Pools and Allocation Limits \[page 272\]](#)

[SAP HANA Memory Usage and the Operating System \[page 273\]](#)

[Managing and Monitoring the Performance of SAP HANA \[page 394\]](#)

[Workload Management \[page 621\]](#)

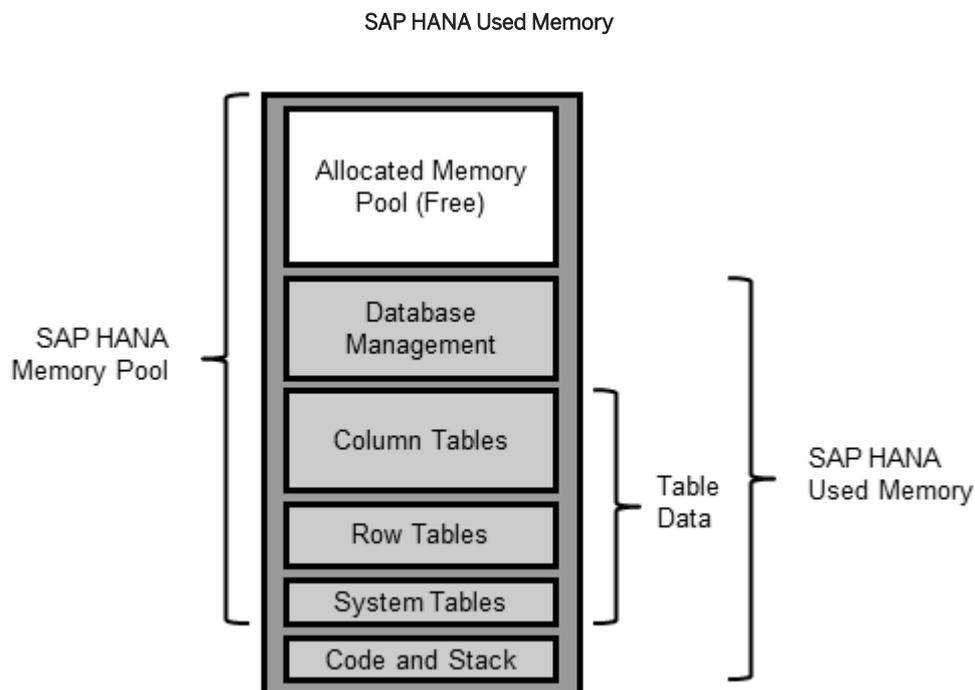
[SAP Note 170449](#)

6.7.5.1 SAP HANA Used Memory

The total amount of memory used by SAP HANA is referred to as used memory. It includes program code and stack, all data and system tables, and the memory required for temporary computations.

SAP HANA consists of a number of processes running in the Linux operating environment. Under Linux, the operating system (OS) is responsible for reserving memory to all processes. When SAP HANA starts up, the OS reserves memory for the program code (sometimes called the text), the program stack, and static data. It then dynamically reserves additional data memory when requested by the SAP HANA memory manager. Dynamically allocated memory consists of heap memory and shared memory.

The following figure shows used memory, consisting of code, stack, and table data:



Since the code and program stack size are about 6 GB, almost all of used memory is used for table storage, computations, and database management.

Service Used Memory

An SAP HANA system consists of multiple services that all consume memory, in particular the `indexserver` service, the main database service. The index server holds all the data tables and temporary results, and therefore dominates SAP HANA used memory.

Peak Used Memory

Ultimately, it is more important to understand the behavior of used memory over time and under peak loads. For this purpose, SAP HANA has a special used memory indicator called peak used memory. As the value for used memory is a current measurement, peak used memory allows you to keep track of the maximum value for used memory over time.

You can also reset peak used memory. This can be useful if you want to establish the impact of a certain workload on memory usage. So for example, you can reset peak used memory, run the workload, and then examine the new peak used memory value.

Memory Usage of Tables

The dominant part of the used memory in the SAP HANA database is the space used by data tables. Separate measurements are available for column-store tables and row-store tables.

i Note

The SAP HANA database loads column-store tables into memory column by column only upon use. This is sometimes called "lazy loading". This means that columns that are never used will not be loaded and memory waste is avoided. When the SAP HANA database runs out of allocatable memory, it will try to free up some memory by unloading unimportant data (such as caches) and even table columns that have not been used recently. Therefore, if it is important to measure precisely the total, or worst-case, amount of memory used for a particular table, it is important to ensure that the table is first fully loaded into memory. You can do this by loading the table into memory.

Memory Usage of Expensive Statements

Every query and statement consumes memory, for the evaluation of the statement plan, caching, and, mainly the calculation of intermediate and final results. While many statement executions use only a moderate amount of memory, some queries, for instance using unfiltered cross joins, will tax even very large systems.

Expensive statements are individual SQL statements whose execution time exceeded a configured threshold. The expensive statements trace records information about these statements for further analysis. If in addition to activating the expensive statements trace, you enable per-statement memory tracking, the expensive statements trace will also show the peak memory size used to execute expensive statements.

It is further possible to protect an SAP HANA system against excessive memory usage due to uncontrolled queries by limiting the amount of memory used by single statement executions per host.

Related Information

[Monitoring and Analyzing with the Performance Monitor \[page 397\]](#)

[Monitor Tables by Size and Usage \[page 334\]](#)

[Load/Unload a Column Table into/from Memory \[page 520\]](#)

[Monitoring and Analyzing with the Statements Monitor \[page 406\]](#)

[Monitoring and Analyzing Expensive Statements \[page 407\]](#)

[Setting a Memory Limit for SQL Statements \[page 633\]](#)

6.7.5.2 Memory Sizing

Memory sizing is the process of estimating in advance the amount of memory that will be required to run a certain workload on an SAP HANA database. To understand memory sizing, several questions need to be answered.

- What is the size of the data tables that will be stored in the SAP HANA database?
You may be able to estimate this based on the size of your existing data, but unless you precisely know the compression ratio of the existing data and the anticipated growth factor, this estimate may not be accurate.
- What is the expected compression ratio that SAP HANA will apply to these tables?
The column store of the SAP HANA database automatically uses a combination of various advanced compression algorithms (dictionary, RLE, sparse, and so on) to compress each table column separately. The achieved compression ratio depends on many factors, such as the nature of the data, its organization and data types, the presence of repeated values, the number of indexes (SAP HANA requires fewer indexes), and so on.
- How much extra working memory will be required for temporary computations?
The amount of extra memory will depend on the size of the tables (larger tables will create larger intermediate result tables in operations such as joins), but even more on the expected workload in terms of the concurrency and complexity of analytical queries (each concurrent query needs its own workspace).

The following SAP Notes provide additional tools and information to help you size the required amount of memory:

- SAP Note 1514966 - SAP HANA 1.0: Sizing SAP In-Memory Database
- SAP Note 1637145 - SAP BW on HANA: Sizing SAP In-Memory Database
- SAP Note 2296290 - New Sizing Report for BW on HANA

However, the most accurate method is to import several representative tables into an SAP HANA system, measure the memory requirements, and extrapolate from the results.

Related Information

[SAP Note 1514966](#)

[SAP Note 1637145](#)

[SAP Note 2296290](#)

6.7.5.3 Allocated Memory Pools and Allocation Limits

SAP HANA, across its different processes, reserves a pool of memory before actual use. This pool of allocated memory is preallocated from the operating system over time, up to a predefined global allocation limit, and is then efficiently used by SAP HANA as needed.

SAP HANA preallocates and manages its own memory pool, used for storing in-memory table data, thread stacks, temporary results, and other system data structures. When more memory is required for table growth or temporary computations, the SAP HANA memory manager obtains it from the pool. When the pool cannot satisfy the request, the memory manager increases the pool size by requesting more memory from the operating system, up to a predefined allocation limit.

By default, the allocation limit is calculated as follows: 90% of the first 64 GB of available physical memory on the host plus 97% of each further GB.

There is normally no reason to change the value of this variable, unless you purchased a license for less than the total amount of physical memory. In this case, you need to change the global allocation limit to remain in compliance with the license.

❖ Example

- You have a server with 512GB, but purchased an SAP HANA license for only 384 GB. You therefore set the `global_allocation_limit` to 393216 (384 * 1024 MB).
- You have a distributed HANA system on four hosts with 512 GB each, but purchased an SAP HANA license for only 768 GB. Set the `global_allocation_limit` to 196608 (192 * 1024 MB on each host).

Another case in which you may want to limit the size of the memory pool is on development systems with more than one SAP HANA system installed on a single host. This will avoid resource contentions or conflicts.

Service Allocation Limit

In addition to the global allocation limit, each service running on the host has an allocation limit, the service allocation limit. Given that collectively, all services cannot consume more memory than the global allocation limit, each service has what is called an effective allocation limit. The effective allocation limit of a service specifies how much physical memory a service can in reality consume given the current memory consumption of other services.

❖ Example

A single-host system has 100 GB physical memory. Both the global allocation limit and the individual service allocation limits are 92.5% (default values). This means the following:

- Collectively, all services of the SAP HANA database can use a maximum of 92.5 GB.
- Individually, each service can use a maximum of 92.5 GB.

Therefore, if 2 services are running and the current memory pool of service 1 is 50 GB, then the effective allocation limit of service 2 is 42.5 GB. This is because service 1 is already using 50 GB and together they cannot exceed the global allocation limit of 92.5 GB.

What happens when the allocation limit is reached?

Memory is a finite resource. Once the allocation limit has been reached and the pool is exhausted, the memory manager can no longer allocate memory for internal operations without first giving up something else. Buffers and caches are released, and column store tables are unloaded, column by column, based on a least-recently-used order, up to a preset lower limit. When tables are partitioned over several hosts, this is managed on a host-by-host basis; that is, column partitions are unloaded only on hosts with an acute memory shortage.

Table (column or partition) unloading is generally not a good situation since it leads to performance degradation later when the data will have to be reloaded for queries that need them. You can identify pool exhaustion by examining the `M_CS_UNLOADS` system view.

However, it is still possible that the memory manager needs more memory than is available. For example, when too many concurrent transactions use up all memory, or when a particularly complex query performs a cross join on very large tables and creates a huge intermediate result that exceeds the available memory. Such situations can potentially lead to an out-of-memory failure.

Related Information

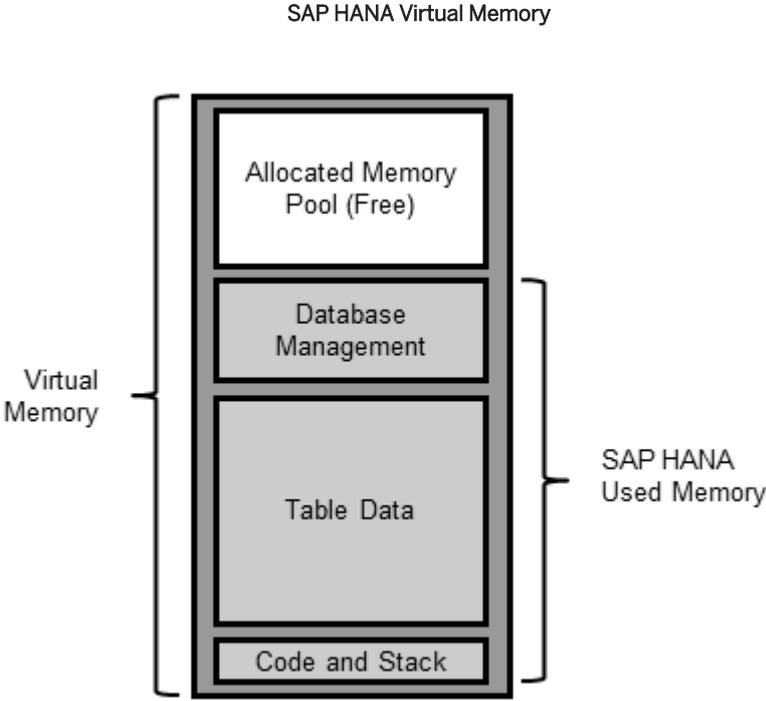
[Change the Global Memory Allocation Limit \[page 478\]](#)

6.7.5.4 SAP HANA Memory Usage and the Operating System

Due to the way in which SAP HANA manages memory, the relationship between Linux memory indicators and SAP HANA's own memory indicators may not correlate as expected.

From the perspective of the Linux operating system, SAP HANA is a collection of separate processes. Linux programs reserve memory for their use from the Linux operating system. The entire reserved memory footprint of a program is referred to as its virtual memory. Each Linux process has its own virtual memory, which grows when the process requests more memory from the operating system, and shrinks when the process relinquishes unused memory. You can think of virtual memory size as the memory amount that the

process has requested (or allocated) from the operating system, including reservations for its code, stack, data, and memory pools under program control. SAP HANA's virtual memory is logically shown in the following figure:



i Note
SAP HANA really consists of several separate processes, so the figure above shows all SAP HANA processes combined.

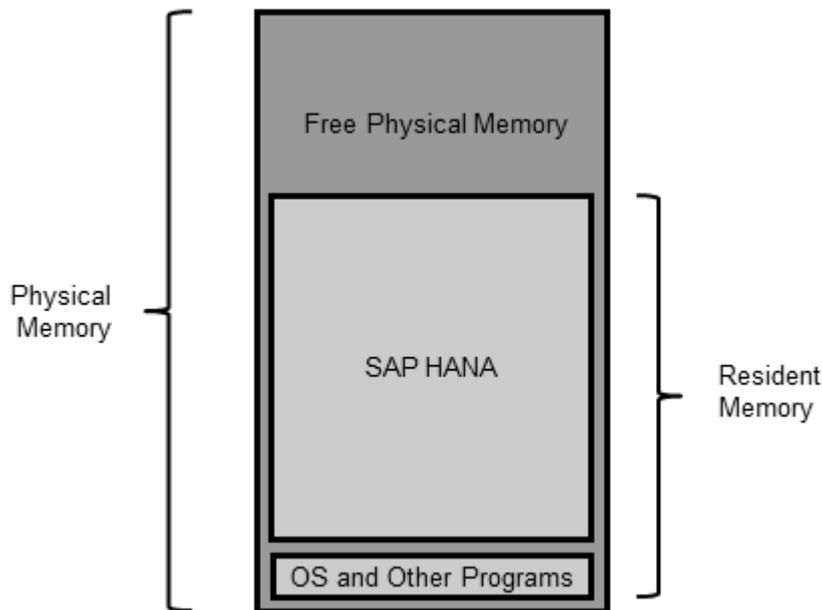
Virtual, Physical, and Resident Memory

When part of the virtually allocated memory actually needs to be used, it is loaded or mapped to the real, physical memory of the host and becomes resident. Physical memory is the DRAM memory installed on the host. On most SAP HANA hosts, it ranges from 256 gigabytes (GB) to 1 terabyte (TB). It is used to run the Linux operating system, SAP HANA, and all other programs.

Resident memory is the physical memory actually in operational use by a process. Over time, the operating system may swap out some of a process's resident memory according to a least-recently-used algorithm to make room for other code or data. Thus, a process's resident memory size may fluctuate independently of its virtual memory size. In a properly-sized SAP HANA appliance, there is enough physical memory, so that swapping is disabled and should not be observed.

This can be illustrated as follows:

SAP HANA Resident Memory



On a typical SAP HANA appliance, the resident memory part of the operating system and all other running programs usually does not exceed 2 GB. The rest of the memory is therefore dedicated for the use of SAP HANA.

When memory is required for table growth or for temporary computations, the SAP HANA code obtains it from the existing memory pool. When the pool cannot satisfy the request, the SAP HANA memory manager will request and reserve more memory from the operating system. At this point, the virtual memory size of SAP HANA processes grows.

Once a temporary computation completes or a table is dropped, the freed memory is returned to the memory manager, which recycles it to its pool without informing the operating system. Therefore, from SAP HANA's perspective, the amount of used memory shrinks, but the processes' virtual and resident memory sizes are not affected. This creates a situation where the used memory value may shrink to below the size of SAP HANA's resident memory. This is normal.

i Note

The memory manager may also choose to return memory back to the operating system, for example when the pool is close to the allocation limit and contains large unused parts.

Related Information

[SAP HANA Used Memory \[page 269\]](#)

[Memory Sizing \[page 271\]](#)

[Allocated Memory Pools and Allocation Limits \[page 272\]](#)

6.7.5.5 Change the Global Memory Allocation Limit

The SAP HANA database preallocates a pool of memory from the operating system over time, up to a predefined global allocation limit. You can change the default global allocation limit.

Prerequisites

You have the system privilege INIFILE ADMIN.

Context

The `global_allocation_limit` parameter is used to limit the amount of memory that can be used by the database. The unit for this parameter is MB. The default value is 0 in which case the global allocation limit is calculated as follows: 90% of the first 64 GB of available physical memory on the host plus 97% of each further GB. Or, in the case of small physical memory, physical memory minus 1 GB.

Changing this parameter does not require a restart.

Procedure

In the `global.ini` configuration file change the value of the `global_allocation_limit` in the `memorymanager` section.

You can enter a value for the entire system and for individual hosts. If you enter only a value for the system, it is used for all hosts. For example, if you have 5 hosts and you set the limit to 5 GB, the database can use up to 5 GB on each host (25 GB in total). If you enter a value for a specific host, then for that host, the specific value is used and the system value is only used for all other hosts. This is relevant only for multiple-host (distributed) systems.

Refer to the *Workload Management* section for details of other options for managing memory including setting a statement limit and admission control.

Related Information

[Allocated Memory Pools and Allocation Limits \[page 272\]](#)

[Setting a Memory Limit for SQL Statements \[page 633\]](#)

[Managing Peak Load \(Admission Control\) \[page 636\]](#)

6.8 Managing Tables

The SAP HANA database stores data in memory in tables, organized in columns, and partitions, distributed among multiple servers.

In addition to the row and column data storage types, the Document Store for JSON artifacts is fully integrated into the SAP HANA database architecture. The document collections of the Document Store are also classified as a type of table and they can be created, updated and read by SQL. The Document Store allows native operation on JSON, for example, filtering, aggregation, and joining JSON documents with HANA column or row store tables.

The default table type is column-type tables. If the table type in a CREATE TABLE statement is not explicitly stated then the table will automatically be a column table. It is possible to override this behavior by setting the configuration parameter `default_table_type` in the `sql` section of the `indexserver.ini` file. The default table type for temporary tables is row-type.

Related Information

[Columnar and Row-Based Data Storage \[page 479\]](#)

[The JSON Document Store \[page 484\]](#)

[Basic Table Management in SAP HANA Cockpit \[page 486\]](#)

[Basic Table Management in SAP HANA Studio \[page 495\]](#)

[Table and Catalog Consistency Checks \[page 508\]](#)

[Memory Management in the Column Store \[page 518\]](#)

[The Delta Merge Operation \[page 526\]](#)

[Data Compression in the Column Store \[page 538\]](#)

[Data Temperature: Extension Nodes \[page 581\]](#)

[Table Partitioning \[page 542\]](#)

[Table Replication \[page 582\]](#)

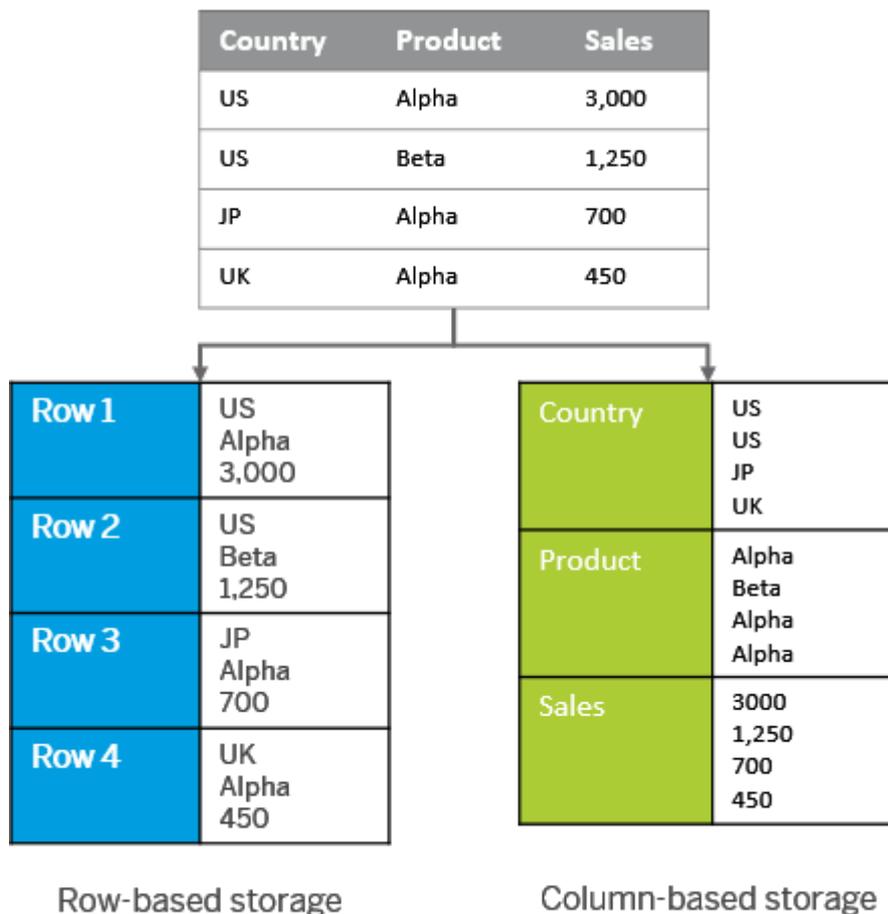
[Table Placement \[page 574\]](#)

[Redistributing Tables in a Scaleout SAP HANA System \[page 600\]](#)

6.8.1 Columnar and Row-Based Data Storage

The SAP HANA database supports two types of table: those that store data either column-wise (column tables) or row-wise (row tables). SAP HANA is optimized for column storage and this is the default table type.

Conceptually, a database table is a two dimensional data structure with cells organized in rows and columns. Computer memory however is organized as a linear sequence. For storing a table in linear memory, two options can be chosen as shown below. A row store stores a sequence of records that contains the fields of one row in the table. In a column store, the entries of a column are stored in contiguous memory locations.



Principle of Row- and Column-Based Storage for a Table

In the SAP HANA database, tables that are organized in columns are optimized for high-performing read operations while still providing good performance for write operations. Efficient data compression is applied to save memory and speed up searches and calculations. Furthermore, some features of the SAP HANA database, such as partitioning, are available only for column tables. Column-based storage is typically suitable for big tables with bulk updates. However, update and insert performance is better on row tables. Row-based storage is typically suitable for small tables with frequent single updates.

The following table outlines the criteria that you can use to decide whether to store your data tables as column tables or row tables:

Storage Type	When to Use
Column store	<ul style="list-style-type: none"> Calculations are typically executed on individual or a small number of columns. The table is searched based on the values of a few columns. The table has a large number of columns. The table has a large number of rows and columnar operations are required (aggregate, scan, and so on) High compression rates can be achieved because the majority of the columns contain only a few distinct values (compared to the number of rows).

Storage Type	When to Use
Row store	<ul style="list-style-type: none"> • The application needs to process only one single record at one time (many selects and /or updates of single records). • The application typically needs to access the complete record. • The columns contain mainly distinct values so compression rate would be low. • Neither aggregations nor fast searching are required. • The table has a small number of rows (for example, configuration tables).

i Note

- The SAP HANA database allows row tables to be joined with column tables. However, it is more efficient to join tables of the same storage type.
- It is possible to change an existing table from one storage type to the other (`ALTER TABLE ALTER TYPE`).

Advantages of Column-Based Storage

Column tables have several advantages:

- **Higher data compression rates**
Columnar data storage allows for highly efficient compression. Especially if the column is sorted, there will be ranges of the same values in contiguous memory, so compression methods such as run length encoding or cluster encoding can be used more effectively.
- **Higher performance for column operations**
With columnar data organization, operations on single columns, such as searching or aggregations, can be implemented as loops over an array stored in contiguous memory locations. Such an operation has high spatial locality and efficiently utilizes the CPU caches.
In addition, highly efficient data compression not only saves memory but also increases speed.
- **Elimination of additional indexes**
In many cases, columnar data storage eliminates the need for additional index structures since storing data in columns already works like having a built-in index for each column: The column-scanning speed of the in-memory column store and the compression mechanisms (especially dictionary compression) already allow read operations with very high performance. In many cases, it will not be required to have additional index structures. Eliminating indexes reduces memory size, can improve write performance, and reduces development efforts. However, this does not mean that indexes are not used at all in SAP HANA. Primary key fields always have an index and it is possible to create additional indexes, if required. In addition, full text indexes are used to support full-text search.
- **Elimination of materialized aggregates**
Thanks to its column-scanning speed, the column store makes it possible to calculate aggregates on large amounts of data on the fly with high performance. This eliminates the need for materialized aggregates in many cases. Eliminating materialized aggregates has several advantages. It simplifies data model and aggregation logic, which makes development and maintenance more efficient; it allows for a higher level of concurrency because write operations do not require exclusive locks for updating aggregated values; and it ensures that the aggregated values are always up-to-date (materialized aggregates are sometimes updated only at scheduled times).
- **Parallelization**

Column-based storage also simplifies parallel execution using multiple processor cores. In a column store data is already vertically partitioned. That means operations on different columns can easily be processed in parallel.

6.8.1.1 History Tables

SAP HANA supports history tables which allow queries on historical data (also known as time-based queries).

History tables are special database tables that only allow inserts. Write operations on history tables do not physically overwrite existing records. Instead, write operations always insert new versions of the data record into the database. The most recent versions in history tables are called current data. All other versions of the same data object contain historical data. Each row in a history table has timestamp-like system attributes that indicate the time period when the record version in this row was the current one. Historical data can be read by requesting the execution of a query against a historical view of the database (`SELECT ... AS OF time`).

Alternatively, you can put a database session in history mode, so that all subsequent queries are processed against the historical view. Currently, SAP HANA supports only column-based history tables.

The history tables in SAP HANA correspond to system-versioned temporary data in SQL. For system-versioned temporal data, the timestamps are automatically set and indicate the so-called transaction time when the data was current. New versions are automatically created by the system during updates.

Time-based queries are read operations against a consistent view of the database that corresponds to a historical point in time. To enable time-based queries, the involved tables must be created as history tables. Only for history tables a history storage is created and filled. Also see the *SAP HANA SQL and System Views Reference*.

Related Information

[CREATE TABLE](#)

6.8.1.2 System-Versioned Tables

System-versioned tables are part of the SQL standard. They support the tracking of changes on column store tables by capturing the validity period of each record.

Specifically in SAP HANA, each record has a *Valid from* and *Valid to* value; these values are maintained by the system so that the mechanism is invisible to the user.

System-versioned tables always consist of two physical tables:

1. The main table of records that are currently valid.
2. A corresponding history table (a one-to-one correspondence) of archived records. A naming convention is not required but may be helpful, for example, to append "_history" to the name of the main table.

i Note

Note that the *history table* in this context is not the same as SAP HANA history tables in versions before HANA 2 SPS 03.

When you execute DML commands (insert, delete, update) on a system-versioned table, any records which are deleted or updated (and where the *Valid from* and *Valid to* columns are populated) are archived in the history table. Internal triggers are executed to manage the archiving process.

Details of system-versioned tables are given in the *SAP HANA SQL and System Views Reference* - some introductory examples are given here.

Creating System-Versioned Tables

You can either create new system-versioned tables (using `CREATE`) or you can modify an existing table with the `ALTER` keyword so that it becomes a system-versioned table. The key syntax elements are shown here using `CREATE`:

```
create column table my_table
(
  <data columns>
  <validfrom_columnname> timestamp not null generated always as row start,
  <validto_columnname> timestamp not null generated always as row end,
  period for system_time (valid_from, valid_to)
)
with system versioning history table my_table_history [[NOT]VALIDATED];
```

- The `valid_from` and `valid_to` columns are timestamps and are maintained by the system; they have a special generation type: *generated always as row start* and *generated always as row end*. The *Valid to* value of the current record is always set to `MAX_TIMESTAMP`.
- These valid from and valid to timestamp columns are then identified as *period for system_time* values.
- The `WITH` clause specifies the corresponding history table for the archived data.
- An optional validation keyword is also supported.

History Tables

The history table has essentially the same structure as the main table but must also meet a number of consistency requirements, specifically: the history table does not have the primary key constraint of the current table, and the `valid_from` and `valid_to` columns are defined without generation clause. If you convert an existing table to a history table you must ensure that all requirements for consistency are met.

Only delete operations on history tables are allowed, this is permitted in order to support housekeeping. Inserts and updates are not allowed and generate a 'feature-not-supported' error. As a workaround to this it would be possible to temporarily drop the system versioning on the table and history table and reapply it later.

Archived records in the history table can still be accessed by referring to their timestamp values. To query a history table one or more timestamp values must be provided with one of the following time operators:

- to select data which was valid at a single point in time you can specify a timestamp with the 'timetravel' operator `AS OF`
- to define time ranges you can use the `FROM - TO` syntax or `BETWEEN - AND`

This example gives basic syntax using `AS OF`.

```
SELECT * FROM <system_versioned_table> FOR SYSTEM_TIME AS OF '<utc-timestamp>'
```

Deleting Data

To permanently delete data from a system-versioned table it is necessary to delete data from both main and history tables in the following sequence:

1. Issue a delete statement referencing the main table - this will delete the data in the main table and move the deleted versions into the history table.
2. Issue a delete statement referencing the history table – this will then delete the data permanently.

System Views

Details of system-versioned tables and their corresponding history table are available in the view `SYS.TEMPORAL_TABLES`.

Other relevant columns in other views are:

- `TEMPORAL_TYPE` in the `SYS.TABLES` view. Possible values are: *Temporal* for the current table, *History* or *NULL* for any other table type.
- `GENERATION_TYPE` in the `SYS.TABLE_COLUMNS` view. Possible values are: *ALWAYS AS ROW START* or *ALWAYS AS ROW END*.

6.8.2 The JSON Document Store

The SAP HANA Document Store is a store for JSON artifacts and allows native operations on JSON including filtering, aggregation, and joining JSON documents with HANA column or row store tables.

JSON documents (JavaScript Object Notation) are stored in so-called collections. The content of a JSON document may be deeply structured but unlike XML it does not have a schema. This means that any valid JSON data may be inserted without first declaring its structure.

Collections appear to users like tables and users can work with them in SQL in a similar fashion. For example, data can be inserted with the regular `INSERT` statement and read via `SELECT`. You can read data from tables and collections in a single statement and you can combine tables and collections by joining as with any other column or row store table.

SAP HANA transactions span all three storage types (Column, Row and Document stores) in a single database which conforms to all the principles of data management: atomicity, consistency, isolation and durability.

Technically, collections are stored in a dedicated binary format; in SQL however, they are defined as tables with their own sub-type. You can see this by selecting collections on the basis of the table type value 'COLLECTION':

```
SELECT * FROM tables WHERE table_type = 'COLLECTION'
```

Full technical details of the Document Store are given in the *SAP HANA Developer Guide for XS Advanced*.

Details of Document Store-specific SQL statements are given in the *SAP HANA SQL Reference and System Views Guide*. The following example is included here to illustrate how a collection of customer details can be created, updated and read. The INSERT statement illustrates the syntax and structure of nested JSON data and the corresponding SELECT statements show the use of path expressions to navigate the content of the JSON document:

```
create collection Customers;
```

```
insert into Customers values('{"name": "Paul",
                             "address": {
                               "street": "Main Street 10",
                               "city": "Heidelberg"
                             }
                           }');
```

```
select * from Customers where "address"."city" = 'Heidelberg';
```

```
select "address"."city", count(*) from Customers group by "address"."city";
```

Because there is no schema to enforce consistent use of data types, you may need to implement your own policy to ensure that data is inserted consistently.

Enabling the JSON Document Store

The JSON Document Store is an optional feature of the SAP HANA database which you have to create for each tenant database. Only a single Document Store process per tenant is possible. Create the Document Store using the ALTER DATABASE statement:

```
ALTER DATABASE <database name> ADD 'docstore'
```

Note that once you have created the Document Store and added collections to it, it can only be removed from the landscape by first removing all collections. The JSON Document Store runs only in the *docstore* server process and collections cannot be moved to any other server (such as the *indexserver*). An option is available for SAP HANA Express Edition users to run the JSON Document Store as an embedded service in the master index server, this option is designed to conserve resources.

The Document Store does not have a pre-determined SQL port, all communication is routed through ordinary index servers.

Diagnostics

You can use database diagnostics (tracing) to analyze the details of Document Store operations. All Document Store trace components are identifiable by the "docstore*" prefix.

Traces which are explicitly designed to trace SQL strings and intermediate results have the suffix "*_data". These traces may contain sensitive data therefore set the trace level of these data tracers only to DEBUG, INFO or WARNING if tracing sensitive information is acceptable. The trace levels ERROR and FATAL will not print any extra sensitive information and are safe to use in this respect. See also *Configure Traces* for more information.

Log File Consolidation

As with row and column store data, activity in the Document Store is logged for data recovery purposes.

Administrators should ensure that the Document Store log file is consolidated from time to time to write a so-called "checkpoint", which improves the re-load and startup performance. This must be done manually or as a scheduled job using the CHECKPOINT statement as described in the *SAP HANA SQL Reference and System Views Guide*.

Related Information

[Configure Traces in SAP HANA Studio \[page 683\]](#)

6.8.3 Basic Table Management in SAP HANA Cockpit

Using SAP HANA database explorer, you can for example execute SQL statements and database procedures, as well as query information about the database and database objects. SAP HANA database explorer is available from the SAP HANA cockpit.

Related Information

[Open the SAP HANA Database Explorer \(SAP HANA Cockpit\) \[page 52\]](#)

6.8.3.1 About the SAP HANA Database Explorer and the SQL Analyzer

Use the SAP HANA database explorer to query information about the database, as well as view information about your database's catalog objects. Use the SAP HANA SQL analyzer to understand and analyze the execution plans of your queries.

The database explorer is integrated into both the SAP Web IDE for SAP HANA and in the SAP HANA cockpit. The database explorer contains features and functions required by both database administrators and developers. For example:

A catalog browser	View the definitions of all types of catalog objects, for example: tables, views, stored procedures, functions, and synonyms. Also, view the content (data) of your tables and views.
An SQL console	Create SQLScript procedures and queries, and then execute them or analyze their performance using the SQL analyzer.
An SQL analyzer	View detailed information on your queries and evaluate potential bottlenecks and optimizations for these queries. The SQL analyzer is accessible from the SQL console, as well as from the plan trace and expensive statement features in the SAP HANA cockpit.
An MDX console	Create and run MDX queries.
An SQL debugger (SAP Web IDE for SAP HANA)	View the call stack, set break points, view and evaluate expressions and variables. This feature is available only for procedures in HDI containers.

Related Information

[Open the SAP HANA Database Explorer \(SAP HANA Cockpit\) \[page 52\]](#)

[Analyzing Statement Performance \[page 440\]](#)

[SAP Note 2373065](#)

6.8.3.2 Create Tables and Views (SQL)

Create a row-store or column-store table by executing the CREATE TABLE or a CREATE VIEW statement.

Prerequisites

To create a table, you must be authorized to create tables in the selected schema.

To create a view, you must be authorized to create views in the selected schema.

Context

A view is a combination or selection of data from tables modeled to serve a particular purpose. Views appear like readable tables, in other words, database operations that read from tables can also be used to read data from views. For example, you can create a view that simply selects some columns from a table, or a view that selects some columns and some rows according to a filter pattern.

i Note

The following procedure describes how to create a simple table and view. Database and application developers use SAP HANA development tools to create database objects such as tables and views as design-time objects in the repository of the SAP HANA database. For more information, see the *SAP HANA Developer Guide*.

Procedure

Choose one of the following options:

Option	Action
Create a table	<p>Execute a CREATE TABLE statement.</p> <p>The following example creates a Table named A that has two INTEGER columns A and B. Column A the primary key.</p> <pre>CREATE TABLE A (A INT PRIMARY KEY, B INT);</pre>
Create a view	<p>Execute a CREATE VIEW statement.</p> <p>The following example creates a view named v that selects all records from table A.</p> <pre>CREATE VIEW v AS SELECT * FROM A;</pre>

6.8.3.3 Export Tables and Other Catalog Objects

Export catalog object definitions from your database by using the SAP HANA database explorer.

Prerequisites

- You must have the SELECT privilege for the catalog objects that you want to export.
- You must have the EXPORT system privilege.

Context

You cannot export a database object larger than 2 GB. If you are exporting a database object larger than 1 GB, then export the object to the SAP HANA computer rather than to the local computer.

Procedure

1. In the catalog browser, right-click the object that you want to export, and then choose *Export Catalog Object*.

Repeat this step to export additional objects.

2. Specify the location to save the export to:

Option	Description
<i>Download the export to the local computer as a single file</i>	<p>Choose this option to save the export as a compressed file to a directory on your local machine.</p> <p>To avoid sending large amounts of data across the network, do not use this option when exporting large amounts of table data. Instead choose to export to the same computer as the SAP HANA server. An export (as well as an import) will fail if the resulting compressed file is larger than 2GB.</p> <p>The contents of the compressed file are not encrypted, so do not use this option if your local machine is not secure.</p>
<i>Save the export to a directory on the SAP HANA computer</i>	<p>Choose this option to save the export to a directory in the database server's file system.</p> <p>If you specify a different directory from the default, then:</p> <ul style="list-style-type: none">○ The directory path must exist.○ The directory path cannot contain symbolic links.○ When the database is part of a distributed system, specify a directory on a shared disk.○ You can specify the <i>backup</i> or the <i>work</i> directory of the database instance, but you cannot specify any another database instance directory. For example, if the database instance is located in the <code>/usr/sap/HDB/HDB00</code> directory, then you could specify: <code>/usr/sap/HDB/HDB00/backup</code> or <code>/usr/sap/HDB/HDB00/work</code> as a directory for the export. <p>The contents of the export are not encrypted.</p>

3. Specify the scope of the export with the following options:

Export option	Description
<i>Include dependencies</i>	Select this option to export the objects as well as any objects that depend on them. For example, if a table has any triggers or indexes associated with it, then their definitions are exported.
<i>Include table data</i>	<p>When exporting tables, select this option to export the data, in addition to the definitions.</p> <p>Clear this option to export only the table object definitions.</p>

Export option	Description
<i>Number of parallel threads</i>	<p>Increasing the number of threads can speed up the export and affect database performance. Consider the following guidelines when choosing the number of threads to use:</p> <ul style="list-style-type: none"> ○ For a view or procedure, use two or more threads, up to the number of dependent objects. ○ For a whole schema, consider using more than ten threads, up to the number of CPU cores in the system. ○ For an SAP BW or SAP ERP system with tens of thousands of tables, using many threads is reasonable (up to 256).
<i>Column table format</i>	<p>Specify the format to use when exporting tables:</p> <ul style="list-style-type: none"> ○ Choose <i>CSV</i> when you are exporting row-store tables or when you must read the contents of the export. ○ Choose <i>Binary</i> to speed up the export time when you are exporting column-store tables.

4. Choose *Export*.

Results

The catalog objects are exported to the specified location. The length of time that the export process takes depends on number of objects being exported and the scope of the export.

Monitor the progress of a running export by using the M_EXPORT_BINARY_STATUS monitoring view.

Next Steps

Import the exported catalog objects into another SAP HANA database by using the database explorer or by executing the IMPORT statement.

Related Information

[Import Tables and Other Catalog Objects \[page 491\]](#)

6.8.3.4 Import Tables and Other Catalog Objects

Import catalog objects into your database by using the SAP HANA database explorer.

Prerequisites

- The database objects being imported must have been exported by using the export feature in the database explorer.
- In the destination database you must have the INSERT/UPDATE, DROP, or CREATE object privileges for the destination catalog objects and you must have the IMPORT system privilege. By default, these privileges are not granted to an HDI container user; they must be granted.
- If you are importing binary files, then those files must have been generated on an SAP HANA system that has the same endianness as the destination system.

Context

You cannot import a database object larger than 2 GB. If you are importing a database object larger than 1 GB, then import the object to the SAP HANA computer rather than to the local computer.

Procedure

1. In the SAP HANA database explorer, right-click the schema that you want to import the objects to, and choose *Import Catalog Object*.
2. Specify the location where the objects to be imported reside in a zipped file. If this file is greater than 1 GB, then it must be imported from the same computer as the SAP HANA server to avoid sending large amounts of data across the network. The import will fail if the file exceeds 2 GB.
3. Specify the scope of the import by choosing the relevant options:

Option	Description
<i>Include dependencies</i>	Dependent objects of the selected objects are also imported. For example, if a table has any triggers or indexes associated with it, then these definitions are imported.
<i>Include table data</i>	Import the table data along with the table definitions. If you deselect this option, then only the table definitions are imported.
<i>Replace existing objects</i>	If objects with the same names already exist in the destination database, then they are overwritten by the objects being imported.
<i>Number of parallel threads</i>	Increasing the number of threads can speed up the import and affect database performance. Consider the following guidelines when choosing the number of threads to use: <ul style="list-style-type: none">○ For a view or procedure, use two or more threads, up to the number of dependent objects.

Option	Description
	<ul style="list-style-type: none"> ○ For a whole schema, consider using more than ten threads, up to the number of CPU cores in the system. ○ For an SAP BW or SAP ERP system with tens of thousands of tables, using many threads is reasonable (up to 256).

4. Click *Import*.

Results

The catalog objects are imported. The length of time that the export process takes depends on the number of objects being imported and the scope of the import.

You can monitor the progress of a running import with the M_IMPORT_BINARY_STATUS monitoring view.

Related Information

[Export Tables and Other Catalog Objects \[page 488\]](#)

6.8.3.5 Open Catalog Objects

Browse your database's catalog using the SAP HANA database explorer.

Prerequisites

You must be a user of the database that you want to explore and you must have the required privileges to view the catalog items.

Context

Some monitoring and problem analysis may require you to examine individual tables and views, for example, system views provided by the SAP HANA database. Use the catalog browser, which is located in the left pane, to find and open these catalog objects.

Procedure

1. In the catalog browser, choose the database that you want to explore.
If your database is not listed in the catalog browser, then click [Add a database to the Database Explorer \(+\)](#) to add the database.
The catalog browser lists the catalog objects, grouped by schema.
2. Choose an object type to view its objects.
For example, choose [Tables](#) to list the tables in the database.
3. Choose an object to view its definition in an editor in the right pane, or right-click the object to choose a different action.
For example, right-click a table and choose [Open Data](#) to view the table's data.

Related Information

[Add SAP HANA Cockpit Resources and Databases to the SAP HANA Database Explorer \[page 53\]](#)

6.8.3.6 Execute SQL Statements

Execute SQL statements and analyze their results by using the SQL console that is included with the SAP HANA database explorer.

Prerequisites

You must have the required privileges in the SAP HANA database to execute your SQL statements.

Context

The database explorer includes an SQL console that does the following:

- Executes batches of statements, separated by semicolons.
- Includes a code-completion feature, as well as a code-formatting feature. Right-click within the SQL console to run these features.
- By default, prepares SQL statements before executing them. (This behavior can be changed in your user preferences.)
- Includes the SQL analyzer to help you analyze the performance of your queries.

If your browser supports saving content to the local storage, then for the duration of your session, the content of your SQL consoles is saved to your browser. (This behavior can be changed in your user preferences.)

Procedure

1. Open an SQL console from your application:

Application	Action
The SAP HANA cockpit	<ol style="list-style-type: none"> 1. In the cockpit, navigate to the <i>Overview</i> page of the database resource that you want to execute the SQL statements on. 2. From the <i>Links</i> tile, click <i>Execute SQL</i>. <p>The database explorer opens in a new tab in your browser. The catalog browser lists your resource and an SQL console opens, which is connected to your resource.</p> <p>If your database is not listed in the catalog browser, then:</p> <ol style="list-style-type: none"> 1. Choose <i>Add a database to the Database Explorer (+)</i> from the catalog browser toolbar to add the database. 2. Choose <i>Open SQL Console</i> () from the global toolbar.
The SAP Web IDE for SAP HANA	<ol style="list-style-type: none"> 1. Open the database explorer by clicking <i>Database Explorer</i> from the left sidebar. 2. Choose a database or HDI container in the catalog browser. 3. Click <i>Open SQL Console</i> () from the global toolbar. <p>An SQL console opens, which is connected to your database.</p>

2. Specify a SQL statement.

For example, the following statement returns users who have the EXPORT or IMPORT system privilege:

```
SELECT * FROM EFFECTIVE_PRIVILEGE GRANTEES WHERE (OBJECT_TYPE = 'SYSTEMPRIVILEGE') AND (PRIVILEGE = 'EXPORT' OR PRIVILEGE='IMPORT');
```

3. Execute the statement by choosing one of the following options:

Option	Action
Execute all statements.	Click the <i>Run</i> icon () from the global toolbar or press F8.
Execute individual statements.	Highlight the statement, and then click the <i>Run</i> icon () from the global toolbar or press F8.
Execute a query and open the SQL analyzer to view information about the query's plan.	<p>Open the <i>Run</i> dropdown list from the global toolbar, and then choose <i>Analyze SQL</i>.</p> <p>A new tab opens to display the query plan for your statement in the SQL analyzer.</p>
Execute the content of the current line.	Open the <i>Run</i> dropdown list from the global toolbar, and then choose <i>Run Line</i> .
Prompt for the values of parameters before executing the statement.	Open the <i>Run</i> dropdown list from the global toolbar, and then choose <i>Prepare Statement</i> .

Results

The *Result* pane appears with the results. Multiple *Result* tabs appear when there is more than one result set. By default, only the first 1000 rows in a result set are retrieved.

Related Information

[Analyzing Statement Performance \[page 440\]](#)

6.8.4 Basic Table Management in SAP HANA Studio

The SAP HANA studio provides several functions for the basic administration and monitoring of tables and views.

Related Information

[Opening Tables and Views \[page 496\]](#)

[Viewing Options for Tables and Views \[page 497\]](#)

[Export Tables and Other Catalog Objects \[page 503\]](#)

[Import Tables and Other Catalog Objects \[page 505\]](#)

[Import ESRI Shapefiles \[page 507\]](#)

[Create a Table in Runtime \[page 500\]](#)

[Create a View in Runtime \[page 502\]](#)

6.8.4.1 Opening Tables and Views

Some monitoring and problem analysis may require you to examine individual tables and views, for example, system views provided by the SAP HANA database. You can open tables and views in the SAP HANA studio in different ways. Several viewing options are available depending on what you want to do.

6.8.4.1.1 Navigate to a Table or View

Navigate to the table or view in the *Systems* view.

Procedure

1. In the *Systems* view, navigate to the table or view you want to open.
2. From the context menu, choose how you want to view the table or view:
 - Definition
 - Content
 - Data preview

i Note

By default, double-clicking the table or view in the *Systems* view opens its definition. You can configure this setting in the preferences of the SAP HANA studio.

Results

The table or view is displayed using the selected viewing option.

Related Information

[Table Definition \[page 498\]](#)

[Table/View Content \[page 500\]](#)

[Data Preview \[page 500\]](#)

6.8.4.1.2 Search for a Table or View

Search for the table or view.

Procedure

1. From the *Systems* view toolbar, choose the  (*Find Table*) button.
2. Enter a search string (at least two characters, case insensitive).
Matching tables and views are displayed immediately.
3. Select the required table or view.
4. Choose whether you want to display the content and/or the definition of the table or view.

Results

The table or view is displayed using the selected viewing option.

Related Information

[Table Definition \[page 498\]](#)

[Table/View Content \[page 500\]](#)

[Data Preview \[page 500\]](#)

6.8.4.2 Viewing Options for Tables and Views

You can open tables and views in different ways in the SAP HANA studio depending on what you want to do.

Open a table or view using one of the following viewing options:

- [Table Definition \[page 498\]](#)
- [Table/View Content \[page 500\]](#)
- [Data Preview \[page 500\]](#)

6.8.4.2.1 Table Definition

The definition view of a table provides you with information about the table's structure and properties (for example, schema, type, column properties, and indexes).

Detailed information relating to the table's memory usage and size is available on the *Runtime Information* sub-tab. This information can be useful in the following cases, for example:

- You want to examine the memory usage of an individual table in detail as part of performance analysis or optimization.
- You want to review the partitioning of a table.

Due to the different memory management concepts for row store and column store tables, the information displayed varies according to table type.

i Note

For views, only the create statement is available.

Runtime Information of Column-Store Tables

For column-store tables, you can review the following information:

- Overall memory usage information for the table, including total size of the table in memory, size of main and delta storages in memory, number of records, and size on disk

i Note

It is not possible to accurately determine the memory consumption of a table from its size on disk. This is because not all data structures that represent a table are stored on disk, they are only created when the table is loaded into memory.

- Detailed memory usage information at the level of partition and individual column

i Note

If the table is not partitioned, the information for the single item on the *Parts* tab is for the table.

The following information may be useful:

Column	Description
Total size	The cumulative in-memory size of all columns and internal structures in the partition, or of the individual column
Main size	The cumulative in-memory size of all columns in the partition in main storage, or of the individual column in main storage
Delta size	The cumulative in-memory size of all columns in the partition in delta storage, or of the individual column in delta storage

Column	Description
Estimated maximum size	The estimated maximum size of the table when loaded into memory, including main and delta storages
Time of last delta merge operation	Time of last delta merge operation
Load status	Partitions can be fully, partially, or not loaded. Individual columns can be either loaded or not.
Main storage compression ratio (%)	The current compression ratio of the column in main storage

i Note

If you want to analyze the compression ratio of a table, it must be fully loaded into memory.

Runtime Information of Row-Store Tables

For row-store tables, you can review the following information:

- Overall memory usage information for the table, including total size of the table in memory, number of records, and size on disk
- Memory usage of fixed and variable parts of the table
Row tables are permanently stored in memory using a linked list of pages. The values displayed here indicate the occupancy level of available pages.

i Note

The memory usage information of column and row store information displayed on the *Runtime Information* tab is retrieved from the following monitoring views:

- M_TABLES
- M_RS_TABLES
- M_CS_TABLES
- M_CS_COLUMNS
- M_CS_PARTITIONS

Related Information

[Memory Sizing \[page 271\]](#)

6.8.4.2.2 Table/View Content

Opening the content view of a table or view executes a SELECT statement on the table/view. The results set shows the actual records in the table/view. This is useful, for example, if you want to view the content of a system view to help you understand what is happening in the database.

When you open the content view, by default, only the first 1,000 rows of the table or view are displayed. You can change this setting in the preferences of the SAP HANA studio under ► [SAP HANA](#) ► [Runtime](#) ► [Catalog](#) ⌵.

To view the full content of a table cell, for example a large object (LOB) value, in the context menu of the cell, choose ► [Export Cell to...](#) ► [Zoom...](#) ⌵

i Note

LOB values are not formatted. Any LOB data that cannot be visualized is not changed.

6.8.4.2.3 Data Preview

Opening the data preview of a table or view allows you to analyze its content in different ways. Similarly to the content view, this is particularly useful for analyzing system views.

♣ Example

You want to check the global memory consumption of the database over the last 30 days.

1. Open the data preview of the table `HOST_RESOURCE_UTILIZATION_STATISTICS` (`_SYS_STATISTICS` schema).
2. Choose the *Data Analysis* tab.
3. Move the column `SERVER_TIMESTAMP` from the *Available Objects* area to the *Labels Axis* area.
4. Move the column `INSTANCE_TOTAL_MEMORY_USED_SIZE` from the *Available Objects* area to the *Values Axis* area.
5. Choose your preferred graphical output.

6.8.4.3 Create a Table in Runtime

A table is a two dimensional data structure with cells organized in rows and columns. Tables can be created as row-store or column-store tables depending on the use case.

Prerequisites

To create a table, you must be authorized to create objects in the selected schema.

Context

i Note

The following procedure describes how to create a simple table in runtime. Database and application developers use SAP HANA development tools to create database objects such as tables as design-time objects in the repository of the SAP HANA database. For more information, see the *SAP HANA Developer Guide*.

Procedure

1. In the *Systems* view, open the catalog and navigate to the schema in which you want to create the new table.
2. In the context menu of the schema in which you want to create the table, choose *New Table*
3. Enter the following information:
 - Table name
 - Table type (column store or row store)
4. Define the columns of your table as follows:
 - a. Enter the name and properties of the first column.
 - b. To add further columns, choose the  button.
5. If necessary, add indexes.
 - a. On the *Indexes* tab, choose the  button.
 - b. Specify the name and the index type (standard index or full-text index).
A full-text index enables full-text search.
 - c. In the lower part of the screen, define the index for the required column(s), together with any other necessary parameters.

i Note

You can create an index for a table any time either by right-clicking the table in the *Systems* view and choosing *New Index*, or opening the table definition for editing.

Indexes are added to the table definition and in the schema's *Indexes* folder.

6. To create the table, choose  (*Create Table*).

Results

The table appears in the *Tables* folder of the relevant schema.

6.8.4.4 Create a View in Runtime

A view is a combination or selection of data from tables modeled to serve a particular purpose.

Prerequisites

You are authorized to create objects in the selected schema and to select data from the tables to be included in the view.

Context

A view is a combination or selection of data from tables modeled to serve a particular purpose. Views appear like readable tables, in other words, database operations that read from tables can also be used to read data from views. For example, you can create a view that simply selects some columns from a table, or a view that selects some columns and some rows according to a filter pattern.

i Note

The following procedure describes how to create a simple SQL view in runtime. Database and application developers use SAP HANA modeling tools to create database objects such as modeled views as design-time objects in the repository of the SAP HANA database. For more information, see the *SAP HANA Developer Guide*.

Procedure

1. In the *Systems* view, open the catalog and navigate to the *Views* folder in the relevant schema.
2. In the context menu, choose *New View*.
The editor for creating a new view opens.
3. Specify the view name.
4. Select the relevant tables by dragging them from the *Systems* view into the editor area, or by choosing the  (*Insert*) button.
5. To create a join, proceed as follows:
 - a. Drag a column from one table to the column of another table.
 - b. Choose the join type in the *Join Order* area.If you define more than one join, you can define the order in which the joins are executed using drag and drop.
6. Drag and drop the columns to be contained in the result set into the *Columns* area.
You can specify additional constraints or create synonyms for column names here.
7. To preview the data, choose *Data Preview* in the context menu of the editor.

8. To show the equivalent SQL statement, choose *Export SQL* in the context menu of the editor.
9. To create the view, choose the  (*Execute*) button.

Results

The view appears in the *Views* folder of the relevant schema.

6.8.4.5 Export Tables and Other Catalog Objects

Using the SAP HANA studio, you can export all catalog objects to a file system and then import them back into another database. This may be necessary, for example, to move data from a test system to a production system, clone your system, or provide the data to SAP Support so they can replicate a scenario.

Prerequisites

- You have SQL object privilege SELECT for the catalog objects in question.
- To browse server directories the SAP HANA studio and server have to be installed on one machine.

Procedure

1. Select the objects you want to export in one of the following ways:

Option	Description
Find objects for export	<ol style="list-style-type: none"> 1. From the main menu, choose File > Export. 2. Select the export destination SAP HANA > Catalog Objects and choose <i>Next</i>. The <i>Export</i> dialog box appears. 3. Search for the objects you want to export and add them to list of objects to be exported on the right.
Select object for export	<ol style="list-style-type: none"> 1. In the <i>Systems</i> view, select the objects that you want to export. 2. From the context menu, choose <i>Export</i>. The <i>Export</i> dialog box appears. The selected objects are displayed on the right of the dialog box. 3. Optional: Search for additional objects to be exported and add them to list of objects to be exported on the right.

2. Choose *Next*.
3. Specify the scope of the export by choosing the relevant options:

Option	Description
Column Table Format	The file format used for export can be either CSV or binary. Exports in binary format are faster and more compact. However, CSV format is better if you need to use the data in a non-SAP HANA system. It also has the advantage of being human readable. Binary is selected by default.

i Note

Column-store tables, procedures, and sequences can be exported in either binary or CSV format. However, row-store tables can be exported only in CSV format. If you trigger an export of catalog objects that includes row-store tables but you select binary as the format, the row-store tables will be exported but in CSV format.

Export Catalog Objects	<p>Using the following options, you can configure the scope of the export:</p> <ul style="list-style-type: none"> ○ Select <i>Including Data</i> to export both object definitions and data are exported. This option is selected by default. If you deselect this option, only object definitions are exported. This may be useful, for example, if want to copy only the table definition in order to create a new table with the same structure. ○ Select <i>Including Dependencies</i> to export dependent objects of selected objects, that is triggers and indexes. This option is selected by default.
-------------------------------	---

4. Specify the location to which the file is to be exported:

Option	Description
Export Catalog Objects on Server	The selected catalog objects are saved to a directory on the database server file system. The default directory is <code>/usr/sap/<SID><instance>/work</code> .

i Note

If you want to specify a different directory in the server's file system, it must already exist and the database must be authorized to access it.

Export Catalog Objects to Current Client	The selected catalog objects are saved in the specified directory on the client file system.
---	--

i Note

The specified directory must be empty. You can specify a directory that does not exist; it will be created when you start the export.

→ Recommendation

For the export of small tables or catalog-only exports, a CSV export to the client file system is appropriate. However, keep the maximum file size of your operating system in mind. A binary export on the server is recommended for large exports (for example, exports over 2 GB).

5. Enter the number of parallel threads to be used for the export.

The more threads you use, the faster the export will be. This does however impact the performance of the database as more threads use more resources.

→ Recommendation

The following guidelines apply:

- For a view or procedure, use two or more threads, up to the number of dependent objects.
- For a whole schema, consider using more than 10 threads, up to the number of CPU cores in the system.

- For a whole SAP BW or SAP ERP system with tens of thousands of tables, using a large number of threads is reasonable (up to 256).

6. Choose *Finish*.

Results

The catalog objects are exported to the specified location. Depending on the number of objects being exported and the scope of the export, this may take some time.

You can monitor the progress of a running export in the monitoring view M_EXPORT_BINARY_STATUS.

Information messages and errors are recorded in the *Error Log* view (▶ [Window](#) ▶ [Show View](#) ▶ [Error Log](#) ▶).

6.8.4.6 Import Tables and Other Catalog Objects

Using the SAP HANA studio, you can import previously exported catalog objects into another database. This may be necessary, for example, to move data from a test system to a production system, clone your system, or provide the data to SAP Support so they can replicate a scenario.

Prerequisites

- Depending on the import, you have the SQL object privilege INSERT/UPDATE, DROP, or CREATE for the catalog objects in question.
- If you are importing binary files, then those files must have been generated on an SAP HANA system that has the same endianness as the destination system.
- To browse server directories, the SAP HANA studio and server have to be installed on one machine.

Context

You can import previously exported catalog objects into an another database. This may be necessary, for example, if you want to move data from a test system to a production system, if you want to clone your system, or if you want to provide the data to SAP Support so they can replicate a certain scenario.

Procedure

1. Select the objects you want to import in one of the following ways:

Option	Description
Find objects for import	<ol style="list-style-type: none"> From the main menu, choose File > Import. Select the import source SAP HANA > Catalog Objects and choose <i>Next</i>. The <i>Import</i> dialog box appears. Specify the location of previously exported objects: <ul style="list-style-type: none"> For objects exported on the database server file system, choose <i>Import Catalog Objects on Server</i> and enter the directory. The default export directory is entered by default. For objects exported to a location on the client file system, choose <i>Import Catalog Objects from Current Client</i> and browse to the directory. Choose <i>Next</i>. Search for the objects that you want to import and add them to list of objects to be imported on the right. You can filter the search results by format type (CSV or binary).
Select objects for import	<ol style="list-style-type: none"> In the <i>Systems</i> view, select the objects that you want to import. From the context menu, choose <i>Import</i>. The <i>Import</i> dialog box appears. Specify the location of previously exported objects: <ul style="list-style-type: none"> For objects exported on the database server file system, choose <i>Import Catalog Objects on Server</i> and enter the directory. The default directory for export is entered by default. For objects exported to a location on the client file system, choose <i>Import Catalog Objects from Current Client</i> and browse to the directory. Choose <i>Next</i>. Optional: Select any additional objects found at the specified export location that you want to import and add them to list of objects to be imported on the right. For objects exported to a location on the client file system, you can filter the search results by format type (CSV or binary).

- Choose *Next*.
- Specify the scope of the import by choosing the relevant options:

Option	Description
Including Data	<p>Object definitions and data are imported.</p> <div style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p>i Note</p> <p>This option is selected by default. However, the export must have included both definition and data.</p> </div> <p>If you deselect this option, only object definitions are imported. For example, you may want to copy only the table definition in order to create a new table with the same structure.</p>
Including Dependencies	<p>Dependent objects of selected objects are also imported, that is triggers and indexes. This option is selected by default.</p>
Replace Existing Catalog Objects	<p>If the objects already exist, they are overwritten.</p>

- Enter the number of parallel threads to be used for the import.

The more threads you use, the faster the import will be. This does however impact the speed of the database as more threads uses more resources.

→ Recommendation

The following guidelines apply:

- For a view or procedure, use two or more threads, up to the number of dependent objects.
- For a whole schema, consider using more than 10 threads, up to the number of CPU cores in the system.
- For a whole SAP BW or SAP ERP system with tens of thousands of tables, using a large number of threads is reasonable (up to 256).

5. Choose *Finish*.

Results

The catalog objects are imported from the specified location. Depending on the number of objects being imported and the scope of the import, this may take some time.

You can monitor the progress of a running import in the monitoring view M_IMPORT_BINARY_STATUS.

Information messages and errors are recorded in the *Error Log* view (▶ [Window](#) ▶ [Show View](#) ▶ [Error Log](#) ▶).

6.8.4.7 Import ESRI Shapefiles

SAP HANA supports the Environmental System Research Institute, Inc. (ESRI) shapefile format. ESRI shapefiles are used to store geometry data and attribute information for the spatial features in a data set. You can import ESRI shapefiles into SAP HANA column-store tables using the SAP HANA studio.

Prerequisites

You have the required privileges for creating and updating objects in the target schema (INSERT/UPDATE or CREATE ANY) or for creating schemas (CREATE SCHEMA).

Context

Spatial data is data that describes the position, shape, and orientation of objects in a defined space. SAP HANA supports spatial data processing, which for example allows application developers to associate spatial information with their data.

The ESRI shapefile format is a popular geospatial vector data format for representing spatial objects in the form of shapefiles (several files that are used together to define the shape). An ESRI shapefile includes at least

three different files: `.shp`, `.shx`, and `.dbf`. The suffix for the main file is `.shp`, the suffix for the index file is `.shx`, and the suffix for the attribute columns is `.dbf`. All files share the same base name and are frequently combined in a single compressed file.

You can import ESRI shapefiles into dedicated column-store tables using the import feature of SAP HANA studio.

Procedure

1. From the main menu, choose **File > Import**.
2. Select import source **SAP HANA > ERSI Shapefiles** and choose *Next*.
3. Select the system that you want to import the shapefiles into and choose *Next*.

Only shapefiles are available for selection.

4. Enter the schema into which you want to import the shape data.

If the schema does not exist and you have the required privileges, the schema is created.

If the table with the same name as the shapefile already exists, you can choose to overwrite it.

5. Enter the number of parallel threads to be used for the import.

The more threads you use, the faster the import will be. This does however impact the speed of the database as more thread uses more resources.

6. To start the import, choose *Finish*.

Results

The shapefiles are imported into column-store tables in the specified schema. Each shapefile corresponds to one table. Depending on the number of shapefiles being imported, this may take some time.

Information messages and errors are recorded in the *Error Log* view (**Window > Show View > Error Log**).

6.8.5 Table and Catalog Consistency Checks

Using stored procedures and commands available in the SAP HANA database, you can perform a range of consistency checks on the database catalog and on database tables.

You are recommended to integrate consistency checks into your routine maintenance schedule so that any problems can be detected as soon as they occur.

Two command line procedures are available to check table consistency and the database catalog:

```
CALL CHECK_TABLE_CONSISTENCY ()
```

```
CALL CHECK_CATALOG ()
```

Optionally, the table consistency check can be scheduled within the embedded statistics service.

For each procedure a list of checking actions is available, for example, CHECK_COLUMN_TABLES, CHECK_ROW_TABLES, CHECK_PARTITIONING_DATA, and so on; these can all be individually activated or omitted from the check as required. For some of these checks a repair option is supported, for example REPAIR_PARTITIONING_DATA. Additional privileges are required for repair actions, these actions must be explicitly specified and must be run separately from check actions. A complete list of all check and repair actions for the two procedures is available by running `GET_CHECK_ACTIONS()`. Details of these commands, configuration options and the statistics features for table consistency checks are given in the sections which follow.

→ Recommendation

Running database checks affects system performance therefore the checks should be run in a timeframe when the system is not at high load. If you are running an active/active (read enabled) system you can run the checks on the read enabled secondary system, or possibly on a system copy.

Related Information

[Table Consistency Check \[page 509\]](#)

[Catalog Consistency Check \[page 512\]](#)

[Configuration Parameters for the Table Consistency Check \[page 513\]](#)

[Active/Active \(Read Enabled\) \[page 1157\]](#)

[The Statistics Service \[page 389\]](#)

6.8.5.1 Table Consistency Check

The table consistency check is a procedure available in the SAP HANA database that performs a range of consistency check actions on database tables. It can be run from the command line or scheduled within the statistics service.

Manual Execution

To execute the procedure manually, you must have the following system privileges:

- CATALOG READ for check actions (or DATA ADMIN)
- DATA ADMIN for repair actions

Input Parameters

To see details of all check actions which relate to table consistency, including a description of what they do, call the procedure `GET_CHECK_ACTIONS`:

```
CALL GET_CHECK_ACTIONS('CHECK_TABLE_CONSISTENCY')
```

Syntax

The syntax of the table consistency check procedure is as follows:

```
CALL CHECK_TABLE_CONSISTENCY ('<check_action1>[,<check_action2>]', '<schema_name>', '<table_name>')
```

This procedure is also available for the Dynamic Tiering option but the syntax and options supported are different. Refer to the *SAP HANA SQL and System Views Reference* for details.

Use the parameter `check_action` to define one or more specific check actions, or enter **CHECK** as the value to execute all default check actions. Use the parameters `schema_name` and `table_name` to define specific schemas and tables to check, or enter **NULL** as the value for these parameters to check all tables in all schemas.

❖ Example

To perform all default check actions on all tables execute:

```
CALL CHECK_TABLE_CONSISTENCY ('CHECK', NULL, NULL)
```

The results returned are the same as listed above when the command is scheduled in the statistics service.

i Note

Some check actions are contained within others and are therefore not explicitly executed when you execute the CHECK action. Repair actions make changes to the data and are excluded from the CHECK action.

Lower case characters and special characters in schema and table names must be enclosed in double quotes. The syntax, for example, for a table named "ABC/abc" in the SYSTEM schema must be as follows:

```
CALL CHECK_TABLE_CONSISTENCY ('CHECK', 'SYSTEM', '"ABC/abc"');
```

Configuration

A set of ini parameters in the `indexserver.ini` file is available to control the command line table consistency check. These include, for example: startup behavior, timeout values, and 'smart' job scheduling parameters to skip large jobs which may severely impact performance. These are described in detail in a separate subsection.

Two SAP Notes on consistency checks are available including an FAQ Note.

Table Consistency Checks in the Statistics Service

You are recommended to schedule the table consistency check so that it runs automatically at regular intervals. The frequency depends on your scenario.

Table consistency checking can be scheduled in the embedded statistics service using collector `_SYS_STATISTICS.Collector_Global_Table_Consistency`. Run-time parameters are maintained as key-value pairs in the `_SYS_STATISTICS.STATISTICS_PROPERTIES` table and the results of the check (details of any errors which are found) are available in the statistics view `GLOBAL_TABLE_CONSISTENCY`. The statistics server also includes a configurable Table Consistency alert (#83) which checks the number of errors and affected tables detected by the consistency check.

The following property values can be defined:

Key	Default Value
internal.table_consistency.check_actions	check_variable_part_sanity, check_data_container, check_variable_part_double_reference_global, check_partitioning, check_replication, check_table_container
internal.table_consistency.target_schema	NULL (all schemas)
internal.table_consistency.target_table	NULL (all tables)
internal.table_consistency.max_duration	0 (no maximum duration)

Note that by default a relatively small number of check actions is carried out to avoid overloading the system, however, the keyword 'check' is available here which acts as a wildcard to execute all check actions. To activate all check actions update this value as shown here.

```
update _SYS_STATISTICS.STATISTICS_PROPERTIES set value = 'check' where key = 'internal.table_consistency.check_actions'
```

The processing sequence for column store tables (row store tables are typically much smaller) is done in strict rotation based on the value of a last check timestamp (oldest first). This ensures that even if the check is canceled before it completes, all tables will eventually be checked over a number of procedure runs.

Result

The results of the automatically executed checks are logged in the view `GLOBAL_TABLE_CONSISTENCY (_SYS_STATISTICS)` with the columns listed here. If no errors are found, the results table is empty.

- SCHEMA_NAME
- TABLE_NAME
- COLUMN_NAME
- PART_ID
- ERROR_CODE
- ERROR_MESSAGE
- SEVERITY

If errors are found you may wish to contact SAP Support to analyze the results and advise on any required action.

Disabling Automatic Checks

You can temporarily disable the statistics collector and alert by executing the following statements:

```
CALL CHECK_TABLE_CONSISTENCY('SET_COLLECTOR_SCHEDULE', 'status', 'inactive')
```

```
CALL CHECK_TABLE_CONSISTENCY('SET_ALERT_SCHEDULE', 'status', 'inactive')
```

You can re-enable the statistics collector and alert by repeating these calls and setting the 'inactive' value to 'idle'.

Related Information

[SAP Note 1977584](#)

6.8.5.2 Catalog Consistency Check

The catalog consistency check can be run from the command line or be scheduled at the operating system level to perform a range of consistency check actions on the database catalog. The frequency with which you do this depends on your scenario.

→ Recommendation

Do not simultaneously run the catalog check and perform DDL operations (for example, dropping users) since this may cause the check to return multiple errors. Either run the catalog check on the system copy or wait until other operations have completed. Only if you continue to receive errors should you contact SAP Support.

Manual Execution

To execute this procedure, you must have the system privilege CATALOG READ (or DATA ADMIN).

The syntax of the table consistency check call is as follows:

```
CALL CHECK_CATALOG  
( '<action>', '<schema_name>', '<object_name>', '<catalog_object_type>' )
```

The `action` parameter specifies the check action(s) to be performed.

To see details of all check actions which relate to catalog consistency, including a description of what they do, call the procedure [GET_CHECK_ACTIONS](#):

```
CALL GET_CHECK_ACTIONS ('CHECK_CATALOG')
```

Use the parameter `action` to define one or more specific check actions, or enter **CHECK** as the value to execute all available actions. Use the parameters `schema_name` and `object_name` to define specific schemas and objects to check, or enter **NULL** as the value for these parameters to check all objects in all schemas.

Specify **NULL** as the value for the parameter `catalog_object_type`. This parameter is not currently effective and is reserved for future use.

❁ Example

To perform all check actions on all objects of all types, execute the statement:

```
CALL CHECK_CATALOG ('CHECK', NULL, NULL, NULL)
```

Object names are case sensitive and this example shows the use of quotation marks to submit a lower-case table name:

```
CALL CHECK_CATALOG ('CHECK', 'SYSTEM', '"mytest"', 'TABLE');
```

Result

If errors are found the procedure returns a set of results with the following columns: SCHEMA, NAME, OBJECT_TYPE, ERROR_CODE, ERROR_MESSAGE.

If errors are found, you may wish to contact SAP Support to analyze the results and advise on the required action.

6.8.5.3 Configuration Parameters for the Table Consistency Check

A set of configuration parameters in the `indexserver.ini` file is available to control the manual table consistency check.

The configuration parameters that control startup behavior are in the `metadata` and `row_engine` sections of the configuration file. Other parameters for timeout values and parameters to skip large jobs which may severely impact performance are available in the `table_consistency_check` section.

[metadata]

Parameter	<code>enable_startup_consistency_check</code>
Short Description	Enable/disable metadata consistency check during SAP HANA startup
Full Description	The metadata consistency check during startup includes CHECK_CATALOG is executed with check action CHECK_OBJECT_REFERENTIAL_INTEGRITY.
Type	Boolean
Change	Offline
Default	True

[row_engine]

Parameter	<code>consistency_check_at_startup</code>
Short Description	Configure row store consistency check during SAP HANA startup.
Full Description	This parameter is used to configure a row store consistency check via CHECK_TABLE_CONSISTENCY during SAP HANA startup. It's a list parameter and the allowed values are 'table', 'page' and 'index', which perform consistency checks on 'table', 'page' and 'index' respectively. This consistency check can be disabled by setting the parameter value to 'none'.
Type	List of strings
Change	Offline

Parameter	consistency_check_at_startup
Default	table,page,index

Parameter	startup_consistency_check_timeout
Short Description	Maximum duration of consistency check at SAP HANA startup
Full Description	This parameter controls the maximum duration of the row store consistency check executed during SAP HANA startup (see <code>consistency_check_at_startup</code>).
Type	Integer
Unit	Second
Change	Offline
Default	600

[table_consistency_check]

Parameter	check_max_concurrency_percent
Short Description	Maximum concurrency for table consistency check
Full Description	This parameter controls the overall CPU and thread resource consumption of CHECK_TABLE_CONSISTENCY, defined as a percentage of <code>max_concurrency</code> in the <code>global.ini</code> file which is the general limit for the number of concurrently running threads.
Type	Integer
Range	1-100
Change	Online
Default	80

Parameter	enable_table_consistency_check_trace
Short Description	Enable/disable table consistency check tracing
Full Description	This parameter controls if the output of CHECK_TABLE_CONSISTENCY is dumped to trace files with the following naming convention: <code><service>_<host>.<port>.table_consistency_check.<timestamp>.trc</code>
Type	Boolean
Change	Online

Parameter	enable_table_consistency_check_trace
Default	True

Parameter	large_job_threshold
Short Description	Threshold for large jobs
Full Description	The parameters <code>large_job_threshold</code> and <code>max_num_large_jobs</code> can be used to make sure that not too many large tables are checked in parallel. Tables exceeding the defined number of rows (default: 100 million) are considered as large tables.
Type	Integer
Unit	Number of rows
Change	Online
Default	100000000

Parameter	max_duration
Short Description	Define the maximum duration of a CHECK_TABLE_CONSISTENCY call
Full Description	This parameter controls the maximum duration of a CHECK_TABLE_CONSISTENCY call. After the specified time the CHECK_TABLE_CONSISTENCY call is implicitly canceled. The default value 0 refers to unlimited duration.
Type	Integer
Unit	Seconds
Change	Online
Default	0

Parameter	max_num_large_jobs
Short Description	Maximum number of large jobs running in parallel
Full Description	The parameters <code>large_job_threshold</code> and <code>max_num_large_jobs</code> can be used to make sure that not too many large tables are checked in parallel. This parameter controls the maximum number of large tables being checked at the same time.
Type	Integer
Change	Online
Default	4

Parameter	max_result_entry
Short Description	Maximum number of generated errors per node
Full Description	Maximum number of generated errors per SAP HANA node; if more errors are found, the following information is provided: <error_count> errors were found in total. <error_count - max_result_entry> errors are suppressed.
Type	Integer
Change	Online
Default	1000000

Parameter	max_result_entry_per_entity
Short Description	Maximum number of generated error per table
Full Description	Maximum number of generated error per table (partition). If more errors are found, the following information is provided: <error_count> errors were found in this table(partition). <error_count - max_result_entry_per_entity> errors are suppressed
Type	Integer
Change	Online
Default	1000

Parameter	remote_check_timeout
Short Description	Maximum wait time for remote checks
Full Description	CHECK_TABLE_CONSISTENCY terminates if checks on remote tables on secondary node aren't finished within the specified time range. The default value of 86400000 milliseconds refers to a maximum duration of one day.
Type	Integer
Unit	Millisecond
Change	Online
Default	86400000

6.8.5.4 Row Store Reorganization

Row store reorganization may be required from time to time to recover data fragments in memory, you can carry out reorganization either online or offline.

Row store reorganization recovers data fragments in memory that result from the way in which memory segments and pages are allocated. During reorganization fragmented pages are moved to other segments and the resultant empty segments are freed.

If reorganization is necessary you can choose to run the process online when the database is running or offline; better results are achieved by running reorganization offline as the database is restarted.

- The online process uses a single SQL system management statement: RECLAIM DATA SPACE.
- The offline process is run by setting the values of a series of 'row_engine' configuration file parameters in the indexserver.ini file and then restarting the database.

Full details of these processes, the prerequisites and the preparation steps are given in SAP Note 1813245 - SAP HANA DB: Row Store Reorganization. Refer also to the SAP HANA SQL and System Views Reference.

Estimating the Benefit of Row Store Reorganization

To estimate the benefit which might be gained from row store reorganization you can run the REORGANIZE_ROWSTORE () procedure. This can be used to assess how much memory might be recovered using both online and offline reorganization methods.

Online Mode

In online mode two options are available, the first applies to the system as a whole:

```
CALL REORGANIZE_ROWSTORE ('ESTIMATE_MEMORY_SAVING', 'ONLINE');
```

The second `online_table` option supports two additional parameters to estimate the memory saving for an individual named schema and table:

```
CALL REORGANIZE_ROWSTORE ('ESTIMATE_MEMORY_SAVING', 'ONLINE_TABLE', 'S1', 'T1');
```

For the online estimates the procedure returns the *estimated minimum memory* and the *estimated maximum memory* in MB which might be saved.

Offline Mode

The syntax for the offline mode is as follows:

```
CALL REORGANIZE_ROWSTORE ('ESTIMATE_MEMORY_SAVING', 'OFFLINE');
```

In this case the procedure returns the *estimated saved memory size in MB*, the *estimated moved page size in MB* and the *estimated moved page count*.

The procedure also includes a documentation parameter: REORGANIZE_ROWSTORE ('HELP').

Related Information

[SAP Note 1813245](#)

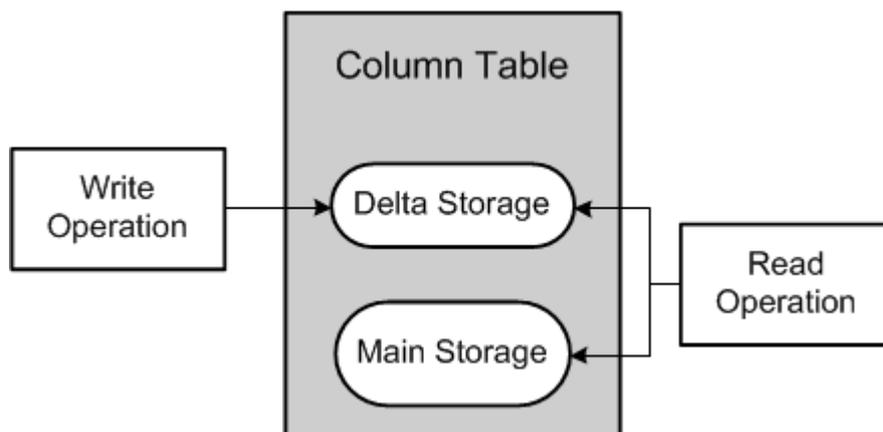
6.8.6 Memory Management in the Column Store

The column store is the part of the SAP HANA database that manages data organized in columns in memory. Tables created as column tables are stored here.

The column store is optimized for read operations but also provides good performance for write operations. This is achieved through 2 data structures: main storage and delta storage.

The main storage contains the main part of the data. Here, efficient data compression is applied to save memory and speed up searches and calculations. Write operations on compressed data in the main storage would however be costly. Therefore, write operations do not directly modify compressed data in the main storage. Instead, all changes are written to a separate data structure called the delta storage. The delta storage uses only basic compression and is optimized for write access. Read operations are performed on both structures, while write operations only affect the delta.

Main Storage and Delta Storage



The purpose of the delta merge operation is to move changes collected in the delta storage to the read-optimized main storage. After the delta merge operation, the content of the main storage is persisted to disk and its compression recalculated and optimized if necessary.

A further result of the delta merge operation is truncation of the delta log. The delta storage structure itself exists only in memory and is not persisted to disk. The column store creates its logical redo log entries for all operations executed on the delta storage. This log is called the delta log. In the event of a system restart, the delta log entries are replayed to rebuild the in-memory delta storages. After the changes in the delta storage have been merged into the main storage, the delta log file is truncated by removing those entries that were written before the merge operation.

i Note

As only data in memory is relevant, the load status of tables is significant. A table can have one of the following load statuses:

- Unloaded, that is, none of the data in the table is loaded to main memory
- Partly loaded, that is, some of the data in the table is loaded to main memory, for example, a few columns recently used in a query
- Fully loaded, that is, all the data in the table is loaded into main memory

However, data that is in the delta storage can only be fully loaded or unloaded. Partial loading is not possible. Therefore, if a delta merge has not been performed and the table's entire data is in the delta storage, the table is either fully loaded or unloaded.

Loading and Unloading of Data in the Column Store

The SAP HANA database aims to keep all relevant data in memory. Standard row tables are loaded into memory when the database is started and remain there as long as it is running. They are not unloaded. Column tables, on the other hand, are loaded on demand, column by column when they are first accessed. This is sometimes called lazy loading. This means that columns that are never used are not loaded and memory waste is avoided.

i Note

This is the **default** behavior of column tables. In the metadata of the table, it is possible to specify that individual columns or the entire table are loaded into memory when the database is started.

The database may actively unload tables or individual columns from memory, for example, if a query or other processes in the database require more memory than is currently available. It does this based on a least recently used algorithm. This unload mechanism can be combined with manually-applied unload priority values. The priority value is applied to a table as a number between 0 and 9; tables with a higher priority are unloaded earlier than other tables.

For more information you can also refer to SAP Note 2127458 *FAQ: SAP HANA Loads and Unloads*

You can also configure columns to allow access to the main storage one page at a time instead of requiring the whole column to be in memory. This enables you to save memory and query a single value in the main storage when certain individual columns or the entire table reside on disk. To enable this feature, specify column description clauses PAGE LOADABLE or COLUMN LOADABLE in the `<column_desc>` of a CREATE TABLE or ALTER TABLE statement.

Related Information

[The Delta Merge Operation \[page 526\]](#)

[SAP Note 2127458](#)

6.8.6.1 Load/Unload a Column Table into/from Memory

Under normal circumstances, the SAP HANA database manages the loading and unloading of tables into and from memory automatically, the aim being to keep all relevant data in memory. However, you can manually load and unload individual tables, as well as load table columns if necessary.

Prerequisites

You have one of the following privileges:

- System privilege TABLE ADMIN
- SQL object privilege UPDATE for the table or the schema in which the table is located

Context

As the SAP HANA database automatically manages the loading and unloading of tables it is not normally necessary to manually load and unload individual tables and table columns. However, this may be necessary for example:

- To precisely measure the total or “worst case” amount of memory used by a particular table (load)
- To actively free up memory (unload)

Load and Unload a Table Using Administration Tools

Context

The procedure is the same in both SAP HANA studio and database explorer using the context menu options available for a selected table in the catalog. In SAP HANA studio you can see detailed information about a table's current memory usage and load status on the *Runtime Information* tab of the table definition. See also monitoring view M_CS_TABLES below.

Procedure

1. In the *Systems* view, navigate to the table in the catalog.
2. In the context menu of the table, choose *Load into Memory* or *Unload from Memory* as required.
3. To start the operation accept the confirmation prompt which is displayed.

Results

If you loaded a table, the complete data of the table, including the data in its delta storage, is loaded into main memory. Depending on the size of the table, this may take some time. The table's load status is FULL.

If you unloaded a table, the complete data of the table, including the data in its delta storage, is unloaded from main memory. Subsequent access to this table will be slower as the data has to be reloaded into memory. The table's load status is NO.

Load and Unload a Table Using SQL

Procedure

1. Open the SQL console and execute the required statement. The LOAD statement supports options to specify one or more table columns, just delta data or all data. You can query the monitoring view M_CS_TABLES for full details of the table (example below).
 - Load MyTable into memory: `LOAD MyTable ALL;`
 - Load columns A and B of MyTable into memory: `LOAD MyTable (A, B);`
 - Unload MyTable from memory: `UNLOAD MyTable;`
2. Query the load status of MyTable: `SELECT loaded FROM m_cs_tables WHERE table_name = 'MyTable';`

Results

If you load only selected columns then the table's load status is PARTIALLY. If you unload a table, the complete data of the table, including the data in its delta storage, is unloaded from main memory.

Related Information

[Memory Sizing \[page 271\]](#)

[Memory Management in the Column Store \[page 518\]](#)

[Table Definition \[page 498\]](#)

6.8.6.2 Managing Memory by Object Usage

You can use the Unused Retention Period feature to automatically unload objects from memory which are not being used.

SAP HANA has a built-in memory management system which automatically unloads swappable resources from memory if the level of available memory gets too low. An additional method for managing memory is the Unused Retention Period, this automatically unloads objects from memory which are not being used.

i Note

You can check the total number of objects and the usage of swappable and non-swappable size using the view `M_MEMORY_OBJECTS`.

Unused Retention Period

To proactively manage memory, even if no low memory situation is present, you can automatically unload swappable objects from memory on the basis of how frequently objects are used. The time-based parameter `unused_retention_period` is available for this in the `global.ini` file.

To use this feature change the default value (initially set to 0) to a number of seconds, for example 7200 (2 hours). Objects which are not used within this time period are flagged as being eligible to unload.

Objects which have exceeded the retention period are not immediately unloaded from memory, an additional checking process which runs at a pre-defined interval initiates the unload. The frequency of the check is controlled by the `unused_retention_period_check_interval` configuration parameter. This is set by default to 7200 seconds (2 hours).

Retention Times and Priorities for Objects

In addition to the general configuration parameter you can apply a retention period and an unload priority value to the table or partition definition itself (see the *SAP HANA SQL and System Views Reference* for details). The following example applies an unload retention period of 60 seconds to table `myTable`.

```
ALTER TABLE "myTable" WITH PARAMETERS ('UNUSED_RETENTION_PERIOD'='60')
```

Note that retention period values can only be applied to tables if a configuration value has been set for `unused_retention_period` in the `global.ini` file.

The unload priority value is a number from 0 to 9 where 0 means the object can never be automatically unloaded and 9 means the earliest unload. The following example reduces the default value of 5 to 2:

```
ALTER TABLE "myTable" UNLOAD PRIORITY 2
```

The unload priority of a table is saved in the `TABLES` table:

```
SELECT UNLOAD_PRIORITY FROM TABLES WHERE TABLE_NAME = "myTable"
```

i Note

Note that changes in unload priority are not immediately effective; the change only takes effect the next time the table is loaded into memory. Therefore if you want to apply the change immediately you should unload and reload the table with the following two statements:

```
UNLOAD "myTable";
```

```
LOAD "myTable" ALL;
```

As these operations may take a long time for large tables, consider running these jobs at a suitable time outside business hours.

Resources have a retention disposition weighting value which also influences the sequence in which objects are unloaded from memory. The weighting is used with the time value since the last access of a resource to calculate a disposition value. Priorities 6-9 correspond to 'early unload' disposition, priorities 1-5 correspond to a 'long term' disposition, and tables with a priority of zero are 'non swappable'. Disposition values are configurable by a set of parameters in the `memoryobjects` section of the `global.ini` file although this is not normally necessary. See also SAP Note 1999997 - FAQ: SAP HANA Memory.

Related Information

[SAP Note 1999997](#) 

6.8.6.3 Hybrid LOBs (Large Objects)

To save memory you can store LOB data on disk, in this case the data is only loaded into memory when it is needed. Alternatively, you can use the configurable Hybrid LOB feature which is flexible and stores LOBs either on disk or in memory depending on their size.

SAP HANA can store large binary objects (LOBs) such as images or videos on disk and not inside column or row structures in main memory. This influences the size of the row store loaded into memory and therefore affects start up and takeover times. An LOB saved on disk is referenced only by an ID in the corresponding table column and is loaded into memory on demand.

This significantly reduces main memory consumption especially when LOB data is not actually requested. LOB data has a short-term disposition setting and if there are memory shortages it is removed from memory before column or table data needs to be unloaded (see Cache Consumption below).

The basic options for managing storage of LOBs are:

- Save all LOBs in memory (LOB size of up to 1GB supported)
- Save all LOBs on disk (a virtual file is created per LOB)
- Use the configurable Hybrid feature which uses three storage types for LOBs depending on their size: the smallest LOBs are stored in memory, the largest LOBs are stored on disk, and medium-sized LOBs are stored together in LOB containers.

Configuration for Hybrid LOBs

Three configuration file settings in the `indexserver.ini` are used to implement and manage the hybrid LOB feature.

The `default_lob_storage_type` setting (section SQL) can be set to `memory` or `hybrid` (default). If Hybrid is activated the following two parameters are also used.

For hybrid LOBs the `lob_memory_threshold` setting (section SQL) must be set to a numeric value of bytes (default 1000). This defines an initial size threshold for the smallest LOBs.

The `midsizelob_threshold` setting (section Persistence) defines the upper threshold for mid-size LOBs and is set to 4000 (bytes) by default.

LOBs are then categorized and stored on the basis of these thresholds:

- Small LOBs (type *Inplace*) are always completely loaded into main memory when the attribute is loaded.
- Medium-sized LOBs where the size is between the two thresholds have the storage type *Packed*. These LOBs share a single LobContainer per attribute and are only loaded into main memory when required. You can prevent mid size LOBs being created by setting `lob_memory_threshold` to the same value as `midsizelob_threshold`.
- Large LOBs (type *File*) are stored in their own virtual file on disk and are loaded individually into main memory on demand.

The Hybrid LOB feature is available in the row and column store but type *Packed* is only implemented for column store tables. If necessary, the feature can be applied to tables retrospectively using the ALTER TABLE command. This may be required for tables that pre-date this feature (see Migrate below).

Example SQL code to apply these settings to a newly-created table is given here applying values for the `lob_memory_threshold` setting. Note that this includes a value for the LOB data type (blob, clob and so on). The effects of applying values of 0 (all LOBs saved on disk) and null (all LOBs saved in memory) are also shown here.

```
create column table <table> (id int, data blob memory threshold 0); -- all lobs
are on disk
create column table <table> (id int, data clob memory threshold 1000); -- all
lobs <= 1000 bytes are in memory, larger lobs are on disk
create column table <table> (id int, data nclob memory threshold null); -- all
lobs are in memory
```

Note that `memory threshold` is always referenced as smaller or equal (`<=`).

The ALTER TABLE command also offers an lob reorganize clause which applies the `midsizelob_threshold` system property to specified columns in a table. See the ALTER TABLE section of Data Definition Statements in the *SAP HANA SQL and System Views Reference* for more information.

Memory Consumption per Table

Even when LOB data is stored on disk you still need to store some information in memory. This information can be retrieved by querying M_TABLES, M_CS_TABLES or M_CS_COLUMNS:

```
SELECT * FROM M_CS_TABLES WHERE SCHEMA_NAME = '<schema>' ORDER BY
MEMORY_SIZE_IN_TOTAL DESC;
SELECT * FROM M_CS_COLUMNS WHERE SCHEMA_NAME = '<schema>' AND TABLE_NAME =
'<table>' ORDER BY MEMORY_SIZE_IN_TOTAL DESC;
```

System HEAP_MEMORY

To see how much main memory is consumed by hybrid LOB data stored on disk that is actually loaded into main memory use M_HEAP_MEMORY:

```
SELECT * FROM M_HEAP_MEMORY WHERE CATEGORY = 'Pool/PersistenceManager/LOBContainerDirectory';
```

Cache Consumption

To speed up LOB data access when they are stored on disk, LOB data is cached inside SAP HANA page cache with short term disposition.

i Note

Do not change this as it will cause performance issues.

During high load HEAP_MEMORY might increase significantly (until SAP HANA's general memory limit is reached). This is no problem as LOB data is unloaded first from the page cache as it uses short term disposition.

For memory analysis the cache may be cleaned or SAP HANA is restarted in order to free caches. Both options should be used carefully as this unloads all tables and reload might be expensive (meaning it may require a downtime).

For an overview of the cache use: `SELECT * FROM M_MEMORY_OBJECT_DISPOSITIONS`

M_MEMORY_OBJECT_DISPOSITIONS shows which component holds which kind of disposition resource (whether the memory objects are short, mid, long-term or non-swappable). It does not tell you what data is stored in cached pages.

Related Information

[M_TABLE_LOB_FILES](#)

[M_MEMORY_OBJECT_DISPOSITIONS](#)

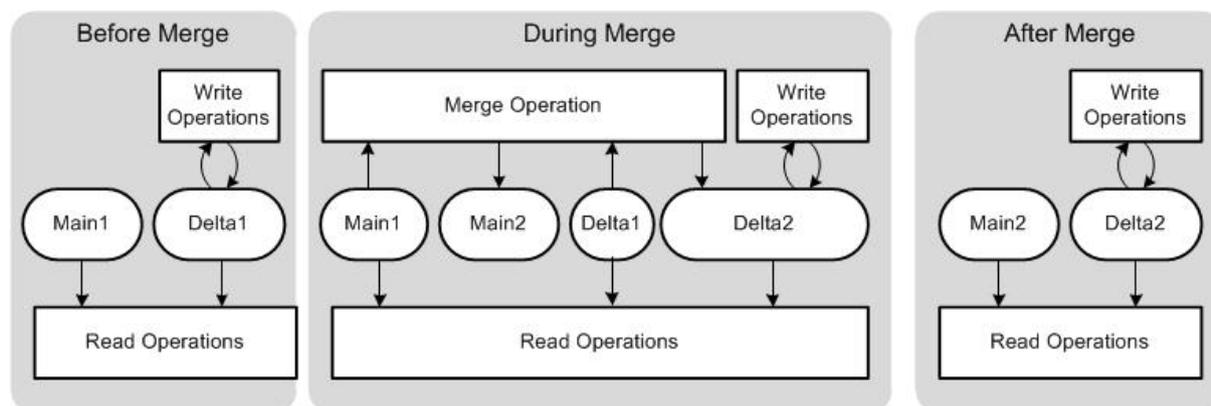
[Data Definition Statements](#)

6.8.7 The Delta Merge Operation

Write operations are only performed on the delta storage. In order to transform the data into a format that is optimized in terms of memory consumption and read performance, it must be transferred to the main storage. This is accomplished by the delta merge operation.

The following figure shows the different steps in the merge process, which objects are involved, and how they are accessed:

The Delta Merge Process



1. Before the merge operation, all write operations go to Delta 1 storage and all read operations read from Main 1 and Delta 1 storages.
2. While the merge operation is running, the following happens:
 1. All write operations go to the second delta storage, Delta 2.
 2. Read operations read from the original main storage, Main 1, and from both delta storages, Delta 1 and Delta 2.
 3. Uncommitted changes in Delta1 are copied to Delta2.
 4. The content of Main 1 and the committed entries in Delta 1 are merged into the new main storage, Main 2.
3. After the merge operation has completed, the following happens:
 1. Main1 and Delta1 storages are deleted.
 2. The compression of the new main storage (Main 2) is reevaluated and optimized. If necessary, this operation reorders rows and adjust compression parameters. If compression has changed, columns are immediately reloaded into memory.
 3. The content of the complete main storage is persisted to disk.

i Note

With this double buffer concept, the table only needs to be locked for a short time: at the beginning of the process when open transactions are moved to Delta2, and at the end of the process when the storages are “switched”.

⚠ Caution

The minimum memory requirement for the delta merge operation includes the current size of main storage plus future size of main storage plus current size of delta storage plus some additional memory. It is important to understand that even if a column store table is unloaded or partly loaded, the whole table is loaded into memory to perform the delta merge.

The delta merge operation can therefore be expensive for the following main reasons:

- The complete main storages of all columns of the table are re-written in memory. This consumes some CPU resources and at least temporarily duplicates the memory needed for the main storages (while Main 1 and Main 2 exist in parallel).
- The complete main storages are persisted to disk, even if only a relatively small number of records were changed. This creates disk I/O load.

This potentially negative impact on performance can be mitigated by the following strategies:

- Executing memory-only merges
A memory-only merge affects only the in-memory structures and does not persist any data.
- Splitting tables
The performance of the delta merge depends on the size of the main storage. This size can be reduced by splitting the table into multiple partitions, each with its own main and delta storages. The delta merge operation is performed at partition level and only for partitions that actually require it. This means that less data needs to be merged and persisted. Note that there are disadvantages to partitioning tables that should also be considered.

Delta Merge on Partitioned Tables

During the delta merge operation, every partition of a partitioned table is treated internally as a standalone table with its own data and delta store. Only the affected partitions are subject to the merge operation. As described above, the whole table has to be duplicated during the merge operation, so for partitioned tables, the amount of needed main memory during the merge operation is reduced, depending on the size of the partition.

⚠ Caution

Before a table is (re-)partitioned, a delta merge operation is executed. Therefore, in the case of huge tables, you have to partition them in good time so as not to run out of memory during the merge operation.

Related Information

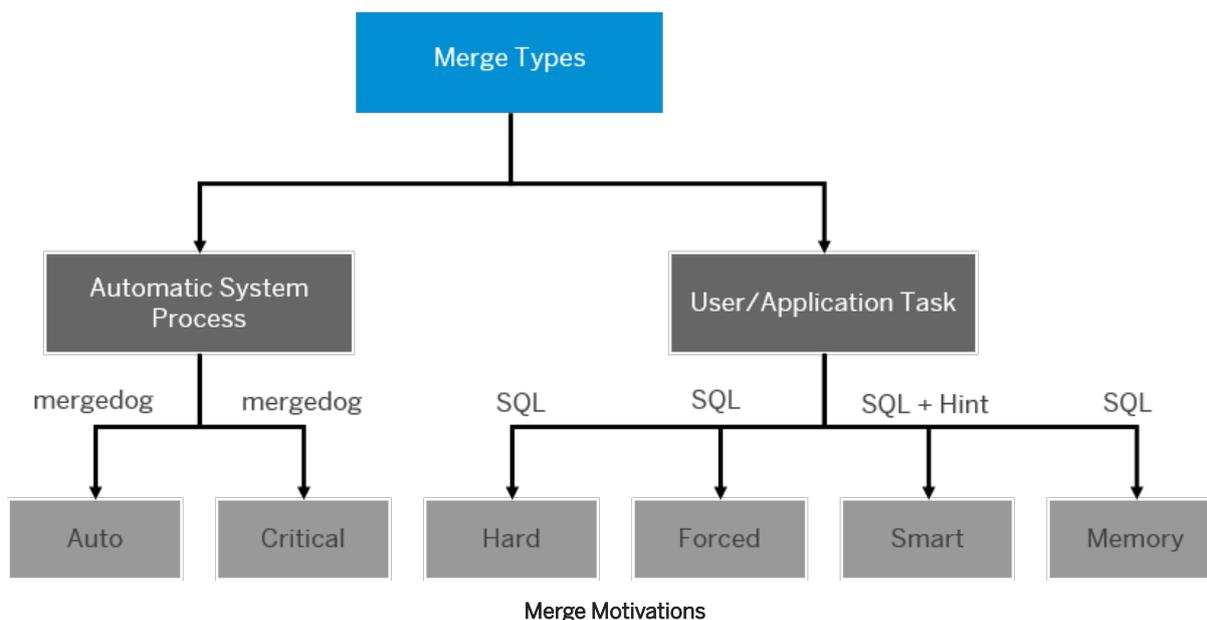
[Table Partitioning \[page 542\]](#)

[Table Partitioning \[page 542\]](#)

6.8.7.1 Merge Motivations

The request to merge the delta storage of a table into its main storage can be triggered in several ways. These are called merge motivations.

The following figure illustrates the different merge motivations and how they are triggered.



Auto Merge

The standard method for initiating a merge in SAP HANA is the auto merge. A system process called `mergedog` periodically checks the column store tables that are loaded locally and determines for each individual table (or single partition of a split table) whether or not a merge is necessary based on configurable criteria (for example, size of delta storage, available memory, time since last merge, and others).

Auto merge is active if the `active` parameter in the `mergedog` section of the `indexserver.ini` file is set to `yes`.

i Note

Auto merge can be activated and deactivated for an individual table in the system view `TABLES (SYS)`. The value in the column `AUTO_MERGE_ON` can be changed to `TRUE` or `FALSE`.

i Note

Deactivating auto merge for a table implicitly disables automatic compression optimization as well (column `AUTO_OPTIMIZE_COMPRESSION_ON` set to `FALSE`). This is the case even if the `AUTO_OPTIMIZE_COMPRESSION_ON` column is set to `TRUE` in the system view `TABLES`. For tables with auto merge disabled, compression optimization must be triggered manually.

Smart Merge

If an application powered by SAP HANA requires more direct control over the merge process, SAP HANA supports a function that enables the application to request the system to check whether or not a delta merge makes sense now. This function is called smart merge. For example, if an application starts loading relatively large data volumes, a delta merge during the load may have a negative impact both on the load performance

and on other system users. Therefore, the application can disable the auto merge for those tables being loaded and send a “hint” to the database to do a merge once the load has completed.

When the application issues a smart merge hint to the database to trigger a merge, the database evaluates the criteria that determine whether or not a merge is necessary. If the criteria are met, the merge is executed. If the criteria are not met, the database takes no further action and only a subsequent hint from the application will trigger another evaluation of the criteria.

Smart merge is active if the `smart_merge_enabled` parameter in the `mergedog` section of the `indexserver.ini` file is set to `yes`.

⚠ Caution

For tables that you want to merge with the smart merge, you should disable the auto merge. Otherwise, the auto merge and smart merge may interfere with each other.

Hard and Forced Merges

You can trigger the delta merge operation for a table manually by executing the SQL statement `MERGE DELTA OF "<table_name>"`. This is called a hard merge and results in the database executing the delta merge for the table either immediately if sufficient system resources are available, or as soon as sufficient system resources become available. The hard merge is therefore subject to the merge token control mechanism.

If you want the merge to take place immediately regardless of system resource availability, you can pass an optional parameter. A forced merge may be useful in a situation where there is a heavy system load, but a small table needs to be merged or if a missed merge of a certain table is negatively impacting system performance. To execute a forced merge, execute the SQL statement `MERGE DELTA OF '<table_name>' WITH PARAMETERS ('FORCED_MERGE' = 'ON')`.

i Note

Unlike system-triggered delta merge operations, all of the manually-executed delta merge operations listed here do not subsequently trigger an optimization of the compression of the table's new main storage. If the table was compressed before the delta merge operation, it remains compressed with the same compression strategy afterward. If it was not compressed before the delta merge operation, it remains uncompressed afterward. After a manual delta merge, you must therefore trigger compression optimization manually.

Critical Merge

The database can trigger a critical merge in order to keep the system stable. For example, in a situation where auto merge has been disabled and no smart merge hints are sent to the system, the size of the delta storage could grow too large for a successful delta merge to be possible. The system initiates a critical merge automatically when a certain threshold is passed.

Related Information

[Perform a Manual Delta Merge Operation \[page 536\]](#)

[Compress a Column Table Manually \[page 541\]](#)

[SAP Note 2057046](#) 

6.8.7.2 The Merge Monitor

The delta merge operation for column tables is a potentially expensive operation and must be managed according to available resources and priority. This is the responsibility of the merge monitor.

The system uses cost functions to decide which table to merge, when, and in which order. There are also cost functions that control how many tables are merged at the same time and how many threads are used to merge a single table.

The merge monitor is responsible for controlling all merge requests for all column tables on a single host. In a distributed system, every index server has its own merge monitor.

All merge requests must acquire a merge token from the merge monitor. A merge token represents an allocation of system resources and "entitles" the merge to actually start. The merge monitor blocks merge requests if there are not enough system resources available or if the same table is already being merged by another thread. This avoids long waits and delays for other threads for inserting or just reading data.

Depending on current system resource consumption, merge motivation, and the evaluation of the various cost functions, the merge monitor lets single requesting merge threads pass and releases waiting threads.

i Note

There is no option or need to disable, stop, or even kill the merge monitor. The merge monitor is not a thread.

6.8.7.3 Cost Functions

The SAP HANA database decides whether or not to execute a requested delta merge and the order in which to execute multiple requests based on configurable merge criteria or cost functions.

Cost functions can be configured depending on the merge motivation, that is whether the merge is being requested by the automatic system process mergedog (auto merge), by a hint from the application (smart merge), by SQL statement (hard merge), and so on.

Cost functions are evaluated in runtime and configured in the mergedog section of the `indexserver.ini` file. The following cost functions are available:

- `auto_merge_decision_func` and `smart_merge_decision_func`
These cost functions determine whether or not a requested delta merge is executed.
- `auto_merge_priority_func` and `smart_merge_priority_func`
These cost functions determine the priority that is assigned to the delta merge request.

- `critical_merge_decision_func`
This cost function determines whether or not a delta merge is executed. It will run a delta merge to avoid situations that could lead to an out of memory or system crash even if other cost functions have been turned off or fail to run.
- `hard_merge_priority_func`
This cost function determines the priority of hard merges.
- `load_balancing_func`
This cost function determines the allocation of system resources to merge processing.

i Note

The decision cost function is evaluated only once for each merge request. In the case of a merge request triggered by a smart merge hint, if the cost function returns a result of false (that is, the system decides that a delta merge is not required), the request is logged but no further evaluation takes place. Only a new hint can potentially initiate a new delta merge.

The following parameters are available for configuring the cost functions. You can use them to build cost functions for all delta merge configurations.

⚠ Caution

It is not recommended that you change the default settings for delta merge unless instructed to do so by SAP Support.

Parameter	Meaning
DMS	Delta memory size [MB] This refers to the size of the table's delta storage.
TMD	Table merge delay [sec] This refers to the time since the last merge of table
MRC	Main row count [million] This refers to the current number of rows in the main storage of the table.
DMR	Deleted main rows [million] This refers to the number of deleted records not in delta storage, but marked as deleted in main storage. Merging makes sense if there are many deleted rows.
DLS	Delta log size [MB]
DCC	Delta cell count [million] This refers to the current number of cells in the delta storage of the table. For example, if the delta storage contains 3 records, each with 4 columns, then the delta cell count is 12.
DRC	Delta row count [million] This refers to the current number of rows in the delta storage of the table.

Parameter	Meaning
QDW	Queuing delay wait [sec] This refers to the time that a merge thread has been waiting for the merge monitor to allocate its merge tokens. This parameter can be useful if you want to implement a first come first served scheduling strategy.
NAME	Table name [string]
SCHEMA	Schema name [string]
CLA	CPU load average [percentage]
LCC	Logical CPU count
THM	Total heap memory [MB]
AHM	Available heap memory, including memory that could be freed [MB]
DUC	Delta uncommitted row count [million] This refers to the number of uncommitted rows in the delta storage of the table.
MMS	Main memory size [MB]
UPT	Index server uptime [sec]
MMU	Main max udiv [million]
OCRC	(Last) optimize compression row count [million]
CRCSOC	Change row count since (last) optimize compression [million]
RP	Table is range partitioned [boolean]
PAL	Process allocation limit [MB]

Cost Functions Examples

Cost Function Configuration	Meaning
<pre>auto_merge_decision_func = DMS>1000 or TMD>3601 or DCC>800 or DMR>0.2*MRC or DLS>5000</pre>	<p>An automatic delta merge of a table is executed if :</p> <ul style="list-style-type: none"> • The size of its delta storage exceeds 1000 MB, or • It has not been merged in over 60 minutes, or • Its delta cell count exceeds 800 million, or • More than 20% of the records in its main storage were deleted, or • The size of its delta log is greater than 5000 MB
<pre>auto_merge_decision_func = DMS > 1000 or DMS > 42 and weekday(now())=6 and secondtime(now())>secondtime('01:00') and secondtime(now())<secondtime('02:00')</pre>	<p>An automatic delta merge of a table is executed if:</p> <ul style="list-style-type: none"> • The size of its delta storage exceeds 1000 MB, unless • It is Saturday between 1.00 and 2.00, in which case it will be merged if delta storage exceeds 42MB <p>Note the week starts with Monday as day 0.</p>

Cost Function Configuration	Meaning
<ul style="list-style-type: none"> <code>smart_merge_decision_func = DMS>1000 or DCC>800 or DLS>5000</code> <code>smart_merge_priority_func = DMS/1000</code> 	<p>A delta merge request of a table triggered by a smart merge hint is executed if:</p> <ul style="list-style-type: none"> The delta storage size exceed 1000 MB, or The delta cell count in the delta storage is greater than 800 million, or The size of the delta log is greater than 5000 MB <p>The system prioritizes smart merge requests based on the size of the delta storage, that is, tables with the bigger deltas are merged first.</p>
<code>hard_merge_priority_func = QDW</code>	Delta merges triggered by hard merge are prioritized only by queuing delay wait, in other words, on a first in first out basis.
<code>hard_merge_priority_function = 1/(7+MMS)</code>	Delta merges triggered by hard merge are prioritized by table size. Smaller tables are merged first, the idea being to free some memory first before bigger tables start merging.

Related Information

[Change a System Property in SAP HANA Studio \[page 301\]](#)

[Data Compression in the Column Store \[page 538\]](#)

[SAP Note 2057046](#) 

6.8.7.4 Merge Tokens

The delta merge operation can create a heavy load on the system. Therefore, controls need to be applied to ensure that merge operations do not consume all system resources. The control mechanism is based on the allocation of merge tokens to each merge operation.

With the exception of the forced merge, a merge operation cannot start unless it has been allocated tokens. If all merge tokens are taken, merge requests have to wait either until the system releases new merge tokens because more resources are available, or until merge tokens have been released by completed merge requests.

The number of merge tokens available for allocation is adjusted based on current system resource availability. This number is recalculated periodically by the system based on a cost function configured in the `load_balancing_func` parameter in the `mergedog` section of the `indexserver.ini` file. The default configuration is `load_balancing_func = 1 + LCC * (100-CLA)/100`. If a hard maximum is required for the amount of tokens available, you can configure a constant value configured or a constant parameter (for example, LCC). Each merge token represents a single CPU.

For every merge request, the number of tokens required to perform the merge is calculated by the system. If the system is not able to determine a value, a default value is returned. This default value can be configured in the `token_per_table` parameter in the `mergedog` section of the `indexserver.ini` file. However, it is not recommended that you change this value.

i Note

It is not possible to check the number of merge tokens available for allocation at any given time, but it is logged in the indexserver trace file if you activate the indexserver component `mergemonitor` with trace level INFO.

Related Information

[Database Trace \(Basic, User-Specific, and End-to-End\) \[page 667\]](#)

6.8.7.5 Monitoring Delta Merge History

Information about all delta merge operations since the last system start are logged in the monitoring view `M_DELTA_MERGE_STATISTICS`. In addition to completed merge operations, information is available on merge hints received by applications and post-merge compression optimization.

You can access a predefined view of these merge statistics in the statement library of the SAP HANA database explorer or in the SAP HANA studio on the [System Information](#) tab of the Administration editor.

The following columns contain potentially useful information:

Column	Description
TYPE	Here you can see the type of merge history entry. The following values are possible: <ul style="list-style-type: none">• MERGE for an actual delta merge operation• HINT for a merge hint sent to SAP HANA by an application• SPARSE for the post-merge optimization of main storage compression
MOTIVATION	This column identifies the underlying merge motivation: AUTO, SMART, HARD, or FORCE
SUCCESS	This column depends on the entry in the TYPE column. <ul style="list-style-type: none">• For MERGE or SPARSE entries, it indicates whether or not the merge or compression optimization operation was successful.• For HINT entries, it indicates whether or not the hint from the application to merge was accepted. If the hint was accepted (SUCCESS=TRUE), then there is an associated entry of type MERGE. If the hint was rejected (SUCCESS=FALSE), then no merge is triggered, so there is no associated MERGE entry.
LAST_ERROR	This column provides information about error codes of the last errors that occurred (most often 2048). Details are provided in ERROR_DESCRIPTION.

i Note

Even if the hint was accepted (SUCCESS=TRUE), this does not necessarily mean that the subsequent merge was successful. You must check the SUCCESS column of the merge entry.

Column	Description
ERROR_DESCRIPTION	<p>The following error codes are possible:</p> <ul style="list-style-type: none"> • Error 2480: The table in question is already being merged. • Error 2481: There are already other smart merge requests for this table in the queue. • Error 2482: The delta storage is empty or the evaluation of the smart merge cost function indicated that a merge is not necessary. • Error 2483: Smart merge is not active (parameter <code>smart_merge_enabled=no</code>). • Error 2484: Memory required to optimize table exceeds heap limit (for failed compression optimization operations (TYPE=SPARSE, SUCCESS=FALSE)).
PASSPORT	For entries with the merge motivation SMART, this column identifies the application that sent the hint to merge (for example, SAP BW powered by SAP HANA)

Note

If the index server is restarted, the delta merge history will initially be empty. The system also collects delta merge statistics in the table `HOST_DELTA_MERGE_STATISTICS (_SYS_STATISTICS)` independent of system restarts. However, as the system only collects statistical data periodically, this table may not have the most recent delta merge operations.

Example

The following is an example of how to use the merge history to find a merge you were expecting to happen based on the settings for triggering smart merge hints in your application.

1. Look for merges triggered by smart merge in the merge history by executing the following SQL statement:

```
SELECT * FROM M_DELTA_MERGE_STATISTICS WHERE table_name = '<your_table>' AND
motivation = 'SMART'
```
2. If no results are returned, check to see if the application actually sent any hints by executing the following statement:

```
SELECT * FROM M_DELTA_MERGE_STATISTICS WHERE type = 'HINT' AND table_name =
'<your_table>'
```

If the application did not send a hint, then the system will not initiate a delta merge. However, if the application did send a hint, the system only executes the merge if the criteria for smart merge are fulfilled. The information is available in the SUCCESS column. The system decides whether or not to accept the hint and execute the merge by evaluating the smart merge decision cost function.
3. If you still have not found the smart merge, check the long-term history by executing the following statement:

```
SELECT * FROM _SYS_STATISTICS.HOST_DELTA_MERGE_STATISTICS WHERE table_name =
'<your_table>'
```

Tracing

You can activate the logging of merge-related information in the database trace for the indexserver component. The relevant trace components are `mergemonitor` and `mergedog`. We recommend the trace level INFO.

Related Information

[View Diagnostic Files in the SAP HANA Database Explorer \[page 662\]](#)

[View Diagnosis Files in SAP HANA Studio \[page 663\]](#)

[Configure Traces in SAP HANA Studio \[page 683\]](#)

[Use the Statement Library to Administer Your Database \[page 356\]](#)

6.8.7.6 Perform a Manual Delta Merge Operation

You can trigger the delta merge operation for a column table manually, for example, if you need to free up memory.

Prerequisites

You have one of the following privileges:

- System privilege TABLE ADMIN
- SQL object privilege UPDATE for the table or the schema in which the table is located

Context

It may be necessary or useful to trigger a merge operation manually in some situations, for example:

- An alert has been issued because a table is exceeding the threshold for the maximum size of delta storage.
- You need to free up memory. Executing a delta merge operation on tables with large delta storages is one strategy for freeing up memory. The delta storage does not compress data well and it may hold old versions of records that are no longer required for consistent reads. For example, you can use the following SQL statement to retrieve the top 100 largest delta storages in memory: `SELECT TOP 100 * from M_CS_TABLES ORDER BY MEMORY_SIZE_IN_DELTA DESC.`

You can trigger the delta merge operation for a column table manually in the SAP HANA studio by menu command or SQL statement. A manually-executed delta merge operation corresponds to a hard merge. However, if you use SQL, you can also pass additional parameters that trigger forced merges and memory-only merges.

Procedure

1. Execute the required merge in one of the following ways:

Option	Description
SAP HANA studio	<ol style="list-style-type: none"> 1. In the <i>Systems</i> view, navigate to the table. 2. In the context menu of the table, choose <i>Perform Delta Merge</i>. 3. Choose <i>OK</i>.
SQL	Open the SQL console and execute the required statement: <ul style="list-style-type: none"> ○ <code>MERGE DELTA OF '<table_name>'</code> (hard merge) ○ <code>MERGE DELTA OF '<table_name>' WITH PARAMETERS ('FORCED_MERGE' = 'ON')</code> (forced merge) ○ <code>MERGE DELTA OF '<table_name>' WITH PARAMETERS ('MEMORY_MERGE' = 'ON')</code> (memory-only merge)

2. Optional: Confirm the delta merge operation in one of the following ways (SAP HANA studio):

- Open the table definition from the *Systems* view and on the *Runtime Information* tab, check the relevant values.

i Note

Even though the delta merge operation moves data from the delta storage to the main storage, the size of the delta storage will not be zero. This could be because while the delta merge operation was taking place, records written by open transactions were moved to the new delta storage. Furthermore, even if the data containers of the delta storage are empty, they still need some space in memory.

- Check the merge history by opening the *Merge Statistics* table on the *System Information* tab. The SUCCESS column indicates whether or not the merge operation was executed.

→ Tip

The delta merge operation can take a long time. You can see the progress of delta merge operations currently running in the Administration editor on the ► *Performance* ► *Job Progress* ▢ tab.

Results

The delta merge operation is executed.

i Note

Unlike system-triggered delta merge operations, manually-executed delta merge operations do not subsequently trigger an optimization of the compression of the table's new main storage. If the table was compressed before the delta merge operation, it remains compressed with the same compression strategy afterward. If it was not compressed before the delta merge operation, it remains uncompressed afterward. After a manual delta merge, you must therefore trigger compression optimization manually.

6.8.8 Data Compression in the Column Store

The column store allows for the efficient compression of data. This makes it less costly for the SAP HANA database to keep data in main memory. It also speeds up searches and calculations.

Data in column tables can have a two-fold compression:

- Dictionary compression
This default method of compression is applied to all columns. It involves the mapping of distinct column values to consecutive numbers, so that instead of the actual value being stored, the typically much smaller consecutive number is stored.
- Advanced compression
Each column can be further compressed using different compression methods, namely prefix encoding, run length encoding (RLE), cluster encoding, sparse encoding, and indirect encoding. The SAP HANA database uses compression algorithms to determine which type of compression is most appropriate for a column. Columns with the PAGE LOADABLE attribute are compressed with the NBit algorithm only.

i Note

Advanced compression is applied only to the main storage of column tables. As the delta storage is optimized for write operations, it has only dictionary compression applied.

Compression is automatically calculated and optimized as part of the delta merge operation. If you create an empty column table, no compression is applied initially as the database cannot know which method is most appropriate. As you start to insert data into the table and the delta merge operation starts being executed at regular intervals, data compression is automatically (re)evaluated and optimized.

Automatic compression optimization is ensured by the parameter `active` in the `optimize_compression` section of the `indexserver.ini` configuration file. This parameter must have the value `yes`.

i Note

If the standard method for initiating a delta merge of the table is disabled (`AUTO_MERGE_ON` column in the system view `TABLES` is set to `FALSE`), automatic compression optimization is implicitly disabled as well. This is the case even if the `AUTO_OPTIMIZE_COMPRESSION_ON` column is set to `TRUE` in the system view `TABLES`. It is necessary to disable auto merge if the delta merge operation of the table is being controlled by a smart merge triggered by the application. For more information, see the section on merge motivations.

Compression Factor

The compression factor refers to the ratio of the uncompressed data size to the compressed data size in SAP HANA.

The uncompressed data volume is a database-independent value that is defined as follows: the nominal record size multiplied by the number of records in the table. The nominal record size is the sum of the sizes of the data types of all columns.

The compressed data volume in SAP HANA is the total size that the table occupies in the main memory of SAP HANA.

❁ Example

You can retrieve this information for a fully-loaded column table from the monitoring view M_CS_TABLES by executing the statement: `select SCHEMA_NAME, TABLE_NAME, MEMORY_SIZE_IN_TOTAL from PUBLIC.M_CS_TABLES where SCHEMA_NAME='<schema>' and TABLE_NAME='<table>'`

The compression factor achieved by the database depends on your SAP HANA implementation and the data involved.

For more information see *Cost Functions*

Cost Functions for Optimize Compression

The cost functions for optimize compression are in the `optimize_compression` section of the service configuration (e.g. `indexserver.ini`)

- **auto_decision_func** - if triggered by MergeDog
- **smart_decision_func** - if triggered by SmartMerge

Default Cost Function Configuration	Meaning
<code>MMU > 0.010240 and if(OCRC, max(MRC, OCRC) / min(MRC, OCRC) >= 1.75, 1) and (not RP or (RP and TMD > 86400))</code>	Optimize compression runs if <ul style="list-style-type: none">• The table contains more than 10240 rows AND• (Optimize compression was never run before OR• The number of rows increase or decrease by factor of 1.75)• AND - if range partitioned, the last delta merge happened more than 24 hours ago.

Related Information

[SAP Note 1514966](#)

[SAP Note 1637145](#)

[Merge Motivations \[page 527\]](#)

[Cost Functions \[page 530\]](#)

[Compress a Column Table Manually \[page 541\]](#)

[Merge Motivations \[page 527\]](#)

6.8.8.1 Check the Compression of a Column Table

For column-store tables, you can check the type of compression applied to table columns, as well as the compression ratio.

Prerequisites

To check the compression status of a table accurately, ensure that it is first fully loaded into main memory.

Procedure

1. To check the type of compression applied to table columns, execute the following SQL statement:

```
SELECT SCHEMA_NAME, TABLE_NAME, COLUMN_NAME, COMPRESSION_TYPE, LOADED from
PUBLIC.M_CS_COLUMNS where SCHEMA_NAME='<your_schema>' and
TABLE_NAME='<your_table>'
```

The columns of the selected table are listed with the type of compression applied. The following values are possible:

- DEFAULT
- SPARSE
- PREFIXED
- CLUSTERED
- INDIRECT
- RLE

i Note

Even if the column is not loaded into memory, the compression type is indicated as DEFAULT. This is because there will always be some level of dictionary compression. However, unless the column is loaded, the database cannot determine the type of compression actually applied. The LOADED column indicates whether or not the column is loaded into memory.

2. Check the compression ratio of table columns, that is, the ratio of the column's uncompressed data size to its compressed data size in memory.

You can do this in the SAP HANA studio:

- a. In the Administration editor, open the table definition in the table editor.
- b. Choose the *Runtime Information* tab.
- c. In the *Details for Table* area, choose the *Columns* tab.

The compression ratio is specified in the *Main Size Compression Ratio [%]* column.

Related Information

[Load/Unload a Column Table into/from Memory \[page 520\]](#)

6.8.8.2 Compress a Column Table Manually

The SAP HANA database decides which columns in a column table to compress and which compression algorithm to apply for each column. It does this as part of the delta merge operation. It is normally not necessary that you interfere with this process. However, you can trigger compression manually.

Prerequisites

You have the UPDATE privilege for the table.

Context

We do not recommend that you interfere with the way in which the SAP HANA database applies compression. However, if a table is not compressed and you think it should be, you can request the database to reevaluate the situation.

Before you do this, consider the reasons why the table may not be compressed, for example:

- The table is very small.
- The table's delta storage has never been merged with its main storage.
- The table was created and filled using an old version of the SAP HANA database that did not compress data automatically. No further data loads, and consequently no delta merge operations, have taken place.
- The auto merge function has been disabled for the table (AUTO_MERGE_ON column in the system view TABLES is set to FALSE). Deactivating auto merge for a columnstore table implicitly disables the automatic compression optimization as well. This is the case even if the AUTO_OPTIMIZE_COMPRESSION_ON column is set to TRUE in the system view TABLES.

Procedure

1. Request the database to reevaluate compression by executing the SQL statement:

```
UPDATE "<your_table>" WITH PARAMETERS ('OPTIMIZE_COMPRESSION'='YES')
```

The database checks all of the table's columns and determines whether or not they need to be compressed, or whether or not existing compression can be optimized. If this is the case, it compresses the data using the most appropriate compression algorithm. However, note the following:

- The database will only reevaluate compression if the contents of the table have changed significantly since the last time compression was evaluated.

- Even if the database does reevaluate the situation, it may determine that compression is not necessary or cannot be optimized and so changes nothing.
2. Check the compression status of the table.
 3. Optional: If compression has not changed, force the database to reevaluate compression by executing the following SQL statement `UPDATE "<your_table>" WITH PARAMETERS ('OPTIMIZE_COMPRESSION' = 'FORCE')`.
The database checks all of the table's columns and determines whether or not they need to be compressed, or whether or not existing compression can be optimized. If this is the case, it compresses the data using the most appropriate compression algorithm. Note that the database may still determine that compression is not necessary or cannot be optimized and so changes nothing.
 4. Check the compression status of the table.

Related Information

[Check the Compression of a Column Table \[page 540\]](#)

[The Delta Merge Operation \[page 526\]](#)

6.8.9 Table Partitioning

The partitioning feature of the SAP HANA database splits column-store tables horizontally into disjunctive sub-tables or partitions. In this way, large tables can be broken down into smaller, more manageable parts. Partitioning is typically used in multiple-host systems, but it may also be beneficial in single-host systems.

Partitioning is transparent for SQL queries and data manipulation language (DML) statements. There are additional data definition statements (DDL) for partitioning itself:

- Create table partitions
- Re-partition tables
- Merge partitions to one table
- Add/delete partitions
- Move partitions to other hosts
- Perform the delta merge operation on certain partitions

When a table is partitioned, the split is done in such a way that each partition contains a different set of rows of the table. There are several alternatives available for specifying how the rows are assigned to the partitions of a table, for example, hash partitioning or partitioning by range.

The following are the typical advantages of partitioning:

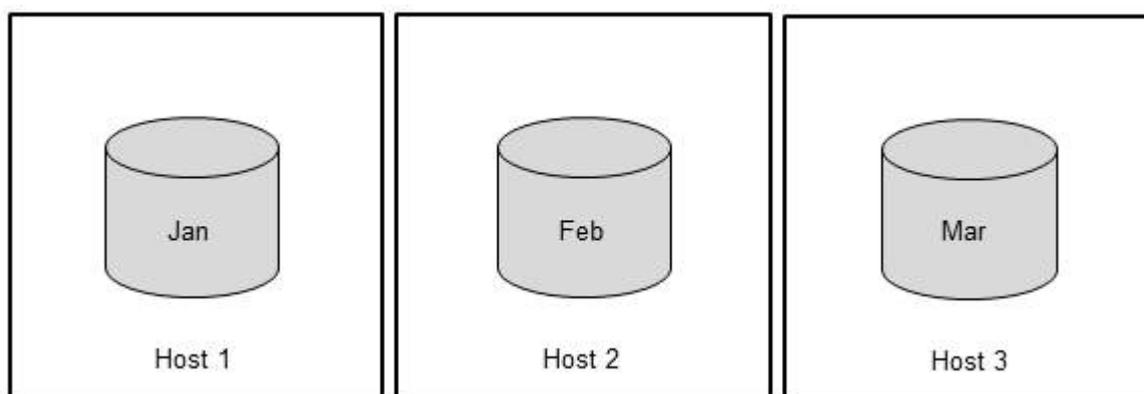
- Load balancing in a distributed system
Individual partitions can be distributed across multiple hosts. This means that a query on a table is not processed by a single server but by all the servers that host partitions.
- Overcoming the size limitation of column-store tables
A non-partitioned table cannot store more than 2 billion rows. It is possible to overcome this limit by distributing the rows across several partitions. Each partition must not contain more than 2 billion rows.
- Parallelization

Partitioning allows operations to be parallelized by using several execution threads for each table.

- Partition pruning
Queries are analyzed to determine whether or not they match the given partitioning specification of a table (static partition pruning) or match the content of specific columns in aging tables (dynamic partition pruning). If a match is found, it is possible to determine the specific partitions that hold the data being queried and avoid accessing and loading into memory partitions which are not required. See *Static and Dynamic Partition Pruning* for details.
- Improved performance of the delta merge operation
The performance of the delta merge operation depends on the size of the main index. If data is only being modified on some partitions, fewer partitions will need to be delta merged and therefore performance will be better.
- Explicit partition handling
Applications may actively control partitions, for example, by adding partitions to store the data for an upcoming month.

The following figure illustrates how a table can be distributed over three hosts with dedicated partitions for individual months.

Example of Table Partitioning



i Note

After adding or removing hosts, it is recommended that you execute a redistribution operation. Based on its configuration, the redistribution operation will suggest a new placement for tables and partitions in the system. If you confirm the redistribution plan, the redistribution operation will re-distribute the tables and partitions accordingly.

For more detailed information about the SQL syntax for partitioning, see *SAP HANA SQL and System Views Reference*.

Related Information

[The Delta Merge Operation \[page 526\]](#)

[Static and Dynamic Partition Pruning \[page 567\]](#)

6.8.9.1 Single-Level Partitioning

When a table is partitioned, its rows are distributed to partitions according to different criteria known as partitioning specifications.

The SAP HANA database supports the following single-level partitioning specifications:

- Round robin
- Hash
- Range

For advanced use cases, these specifications can be nested using multi-level partitioning.

Related Information

[Multi-Level Partitioning \[page 547\]](#)

6.8.9.1.1 Hash Partitioning

Hash partitioning is used to distribute rows to partitions equally for load balancing and to overcome the 2 billion row limitation. The number of the assigned partition is computed by applying a hash function to the value of a specified column. Hash partitioning does not require an in-depth knowledge of the actual content of the table.

For each hash partitioning specification, columns must be specified as partitioning columns. The actual values of these columns are used when the hash value is determined. If the table has a primary key, these partitioning columns must be part of the key. The advantage of this restriction is that a uniqueness check of the key can be performed on the local server. You can use as many partitioning columns as required to achieve a good variety of values for an equal distribution.

For more information about the SQL syntax for partitioning, see *SAP HANA SQL and System Views Reference*.

❁ Example

Creating a Hash-Partitioned Table Using SQL

SQL Command	Result
<pre>CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT, PRIMARY KEY (a,b)) PARTITION BY HASH (a, b) PARTITIONS 4</pre>	<ul style="list-style-type: none">• 4 partitions on columns a and b are created.• The target partition is determined based on the actual values in columns a and b.• At least one column has to be specified.• If a table has a primary key, all partitioning columns must be part of that key.
<pre>CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT, PRIMARY KEY (a,b)) PARTITION BY HASH (a, b) PARTITIONS GET_NUM_SERVERS ()</pre>	The number of partitions is determined by the database at runtime according to its configuration. It is recommended to use this function in scripts, and so on.

Creating Partitions at Specific Locations

You can specify at which location (index server) the partition should be created using the AT LOCATION clause as shown in the following examples. Note that this can be used with all partitioning types: hash, round robin and range.

In this first example, the specific number of partitions to create is given (3) with a corresponding list in the location clause of three hosts (name and port number) so that one partition will be created on each host:

```
CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT) PARTITION BY HASH (a, b)
PARTITIONS 3 AT LOCATION 'myHost01:30001', 'myHost02:30003', 'myHost03:30003';
```

If the number of partitions doesn't match with the number of hosts a best fit is applied automatically: extra locations are ignored if too many are named and locations are reused in round robin fashion if there are more partitions than locations. If no location is specified then the partitions are automatically assigned to hosts at random.

In this second example the number of partitions is returned from the GET_NUM_SERVERS() function:

```
CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT) PARTITION BY HASH (a, b)
PARTITIONS GET_NUM_SERVERS() AT LOCATION 'myHost01:30001';
```

Similarly, in this case, if the number of partitions doesn't match with the number of hosts a best fit is applied automatically.

6.8.9.1.2 Round-Robin Partitioning

Round-robin partitioning is used to achieve an equal distribution of rows to partitions. However, unlike hash partitioning, you do not have to specify partitioning columns. With round-robin partitioning, new rows are assigned to partitions on a rotation basis. The table must not have primary keys.

Hash partitioning is usually more beneficial than round-robin partitioning for the following reasons:

- The partitioning columns cannot be evaluated in a pruning step. Therefore, all partitions are considered in searches and other database operations.
- Depending on the scenario, it is possible that the data within semantically-related tables resides on the same server. Some internal operations may then operate locally instead of retrieving data from a different server.

For more information about the SQL syntax for partitioning, see *SAP HANA SQL and System Views Reference*.

❖ Example

Creating a Round-Robin Partitioned Table Using SQL

SQL Command	Result
<pre>CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT) PARTITION BY ROUNDROBIN PARTITIONS 4</pre>	4 partitions are created. Note: The table must not have primary keys.

SQL Command

```
CREATE COLUMN TABLE MY_TABLE (a INT, b  
INT, c INT) PARTITION BY ROUNDROBIN  
PARTITIONS GET_NUM_SERVERS ()
```

Result

The number of partitions is determined by the database at runtime according to its configuration. It is recommended to use this function in scripts or clients that may operate in various landscapes.

6.8.9.1.3 Range Partitioning

Range partitioning creates dedicated partitions for certain values or value ranges in a table. For example, a range partitioning scheme can be chosen to create a partition for each calendar month. Partitioning requires an in-depth knowledge of the values that are used or are valid for the chosen partitioning column.

Partitions may be created or dropped as needed and applications may choose to use range partitioning to manage data at a fine level of detail, for example, an application may create a partition for an upcoming month so that new data is inserted into that new partition.

i Note

Range partitioning is not well suited for load distribution. Multi-level partitioning specifications address this issue.

The range partitioning specification usually takes ranges of values to determine one partition (the integers 1 to 10 for example) but it is also possible to define a partition for a single value. In this way, a list partitioning known in other database systems can be emulated and combined with range partitioning.

When rows are inserted or modified, the target partition is determined by the defined ranges. If a value does not fit into one of these ranges, an error is raised. To prevent this you can also define an 'others' partition for any values that do not match any of the defined ranges. 'Others' partitions can be created or dropped on-the-fly as required.

Range partitioning is similar to hash partitioning in that the partitioning column must be part of the primary key. Many data types are supported for range partitioning, see the list of data types in *Partitioning Limits* for the complete list.

For more information about the SQL syntax for partitioning, see *SAP HANA SQL and System Views Reference*.

❁ Example

Creating a Range-Partitioned Table Using SQL

The following example creates three columns for integers and divides the first column into four partitions. Ranges are defined using this semantic: `<= VALUES <`, ranges for single values use `=`.

- 1 partition for values greater than or equal to 1 and less than 5
- 1 partition for values greater than or equal to 5 and less than 20
- 1 partition for values of 44
- 1 others partition for all other values which do not match the specified ranges

```
CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT, PRIMARY KEY (a,b))  
PARTITION BY RANGE (a)  
(PARTITION 1 <= VALUES < 5,
```

```
PARTITION 5 <= VALUES < 20,  
PARTITION VALUE = 44, PARTITION OTHERS)
```

The partitioning column (a in this example) has to be part of the primary key.

Data Types

The following data types are allowed for the partitioning column: STRING, TINYINT, SMALLINT, INT, SHORTTEXT, VARCHAR, NVARCHAR, DATE, TIME, TIMESTAMP, SECONDDATE, FIXED, RAW (SQL Binary/Varbinary).

For the RAW datatypes (BINARY and VARBINARY) the partition definition entered in the SQL code must specify the boundaries of the conditions using the HEX representation of the binary value. This is shown in the following example which creates a 2-column table, the first column is for integers and the second column for 16 byte varbinary data. In this case, partition boundaries for the varbinary data must be specified with 32 hexadecimal digits (2 hex characters required for each binary byte):

```
create column table DEMO_RAW (a int, b varbinary(16))  
partition by range (b)  
(  
partition '00000000000000000000000000000000' <= values < '01000000000000000000000000000000',  
partition '01000000000000000000000000000000' <= values < '02000000000000000000000000000000',  
partition '02000000000000000000000000000000' <= values < '03000000000000000000000000000000',  
partition '03000000000000000000000000000000' <= values < 'FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF'  
)
```

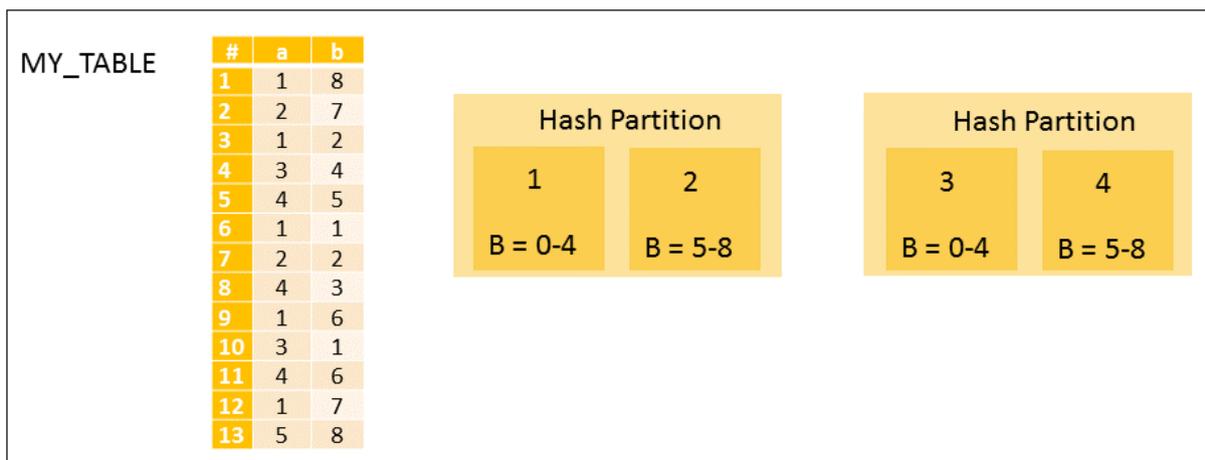
Related Information

[Partitioning Limits \[page 571\]](#)

6.8.9.2 Multi-Level Partitioning

Multi-level partitioning can be used to overcome the limitation of single-level hash partitioning and range partitioning, that is, the limitation of only being able to use key columns as partitioning columns. Multi-level partitioning makes it possible to partition by a column that is not part of the primary key.

The following code example and illustration show how multi-level partitioning can be applied using hash partitioning at the first level and range partitioning at the second level. Data in the second level partitions is grouped on the basis of the value of a selected column 'b': rows where the value is below 5 and rows where the value is 5 or greater but less than 9.



Multi-Level Partitioning

The syntax of the SQL code to create these partitions is as follows:

```
CREATE COLUMN TABLE MY_TABLE (a INT, b INT, PRIMARY KEY (a,b))
PARTITION BY
HASH (a,b) PARTITIONS 2,
RANGE (b) (PARTITION 0 <= VALUES < 5, PARTITION 5 <= VALUES < 9)
```

The primary key column restriction only applies at the first partitioning level. When a row is inserted or updated, the unique constraint of the key must be checked. If the primary key has to be checked on all partitions across the landscape, this would involve expensive remote calls. Second-level partition groups, however, allow inserts to occur whilst only requiring primary key checks on local partitions.

Related second level partitions form a partition group; the figure above shows two groups (partitions 1 and 2 are a group and partitions 3 and 4). When a row is inserted into partition 1, it is only required to check for uniqueness on partitions 1 and 2. All partitions of a partition group must reside on the same host. SQL commands are available to move partitions but it is not possible to move individual partitions of a group, only partition groups as a whole.

Using Date Functions to Partition

You can use multi-level partitioning to implement time-based partitioning to leverage a date column and build partitions according to month or year. This could be used for example to minimize the run-time of the delta merge operation. The performance of the delta merge operation depends on the size of the main index of a table. If data is inserted into a table over time and it also contains a date in its structure, then multi-level partitioning on the date could be very effective. Partitions containing old data are typically only modified infrequently, there is therefore no need for a delta merge on these partitions; the merge is only required on partitions where new data is inserted.

If a table needs to be partitioned by month or by year and it contains only a date column or a timestamp column, you can use the date functions shown below to restrict your query results by year or by year and month.

❖ Example

Partitioning Using Date Functions

This example partitions by hash using the year() function:

```
CREATE COLUMN TABLE MY_TABLE (a DATE, b INT, PRIMARY KEY (a,b)) PARTITION BY HASH  
(year(a)) PARTITIONS 4
```

If a value takes the format "2018-12-08", the hash function is only applied to "2018". This function can also be used for pruning.

This example partitions by range using the year() function:

```
CREATE COLUMN TABLE MY_TABLE (a DATE, b INT, PRIMARY KEY (a,b)) PARTITION BY  
RANGE (year(a)) (PARTITION '2010' <= values < '2013', PARTITION '2013' <= values  
< '2016')
```

This example partitions by range using the year and month value using the month() function:

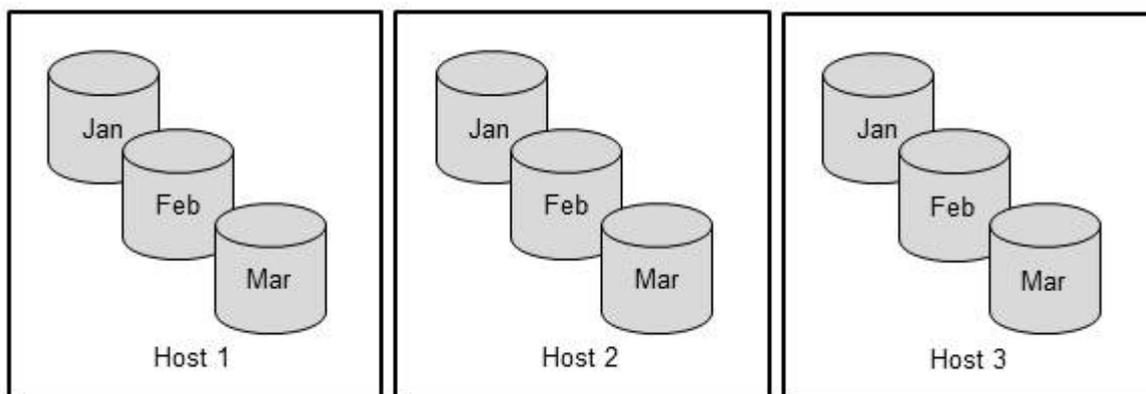
```
CREATE COLUMN TABLE MY_TABLE (a DATE, b INT, PRIMARY KEY (a,b)) PARTITION BY  
RANGE (month(a)) (PARTITION '2005-01' <= values < '2005-07', PARTITION '2005-07'  
<= values < '2006-01')
```

6.8.9.2.1 Hash-Range Multi-Level Partitioning

Hash-range multi-level partitioning is the most common type of multi-level partitioning. Hash partitioning is implemented at the first level for load balancing and range partitioning at the second level for time-based partitioning.

The following figure shows a typical usage scenario. The load is distributed to three hosts using hash partitioning. Range partitioning is used at the second level to distribute the data to individual partitions according to month.

Hash-Range Partitioning



For more information about the SQL syntax for partitioning, see *SAP HANA SQL and System Views Reference*.

❖ Example

Creating a Table with Hash-Range Multi-Level Partitioning Using SQL

```
CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT, PRIMARY KEY (a,b))
PARTITION BY
  HASH (a, b) PARTITIONS 4,
  RANGE (c)
    (PARTITION 1 <= VALUES < 5,
     PARTITION 5 <= VALUES < 20)
```

6.8.9.2.2 Round-Robin-Range Partitioning

Round-robin-range multi-level partitioning is the same as hash-range multi-level partitioning but with round-robin partitioning at the first level.

For more information about the SQL syntax for partitioning, see *SAP HANA SQL and System Views Reference*.

❖ Example

Creating a Table with Round-Robin-Range Partitioning Using SQL

```
CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT)
PARTITION BY
  ROUNDROBIN PARTITIONS 4,
  RANGE (c)
    (PARTITION 1 <= VALUES < 5,
     PARTITION 5 <= VALUES < 20)
```

6.8.9.2.3 Hash-Hash Partitioning

Hash-hash multi-level partitioning is implemented with hash partitioning at both levels. The advantage of this is that the hash partitioning at the second level may be defined on a non-key column.

For more information about the SQL syntax for partitioning, see *SAP HANA SQL and System Views Reference*.

❖ Example

Creating a Table with Hash-Hash Partitioning Using SQL

```
CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT, PRIMARY KEY (a,b))
PARTITION BY
  HASH (a, b) PARTITIONS 4,
  HASH (c) PARTITIONS 7
```

6.8.9.2.4 Range-Range Partitioning

Range-range multi-level partitioning is implemented with range partitioning at both levels. The advantage of this is that the range partitioning at the second level may be defined on a non-key column.

For more information about the SQL syntax for partitioning, see *SAP HANA SQL and System Views Reference*.

❖ Example

Creating a Table with Range-Range Partitioning Using SQL

```
CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT, PRIMARY KEY (a,b))
  PARTITION BY
    RANGE (a)
      (PARTITION 1 <= VALUES < 5,
       PARTITION 5 <= VALUES < 20),
    RANGE (c)
      (PARTITION 1 <= VALUES < 5,
       PARTITION 5 <= VALUES < 20)
```

6.8.9.3 Heterogeneous Partitioning

Heterogeneous partitioning offers more flexibility for range-range two-level partitioning schemas.

The normal second-level range partitioning schema applies the same second level specification to all first level partitions. But for some scenarios a flexible *heterogeneous* second-level range partitioning schema may be required with different second-level range specifications for each first level partition. This is possible for both column-store tables and Dynamic Tiering multi-store tables using the special heterogeneous partitioning clauses for the CREATE TABLE and ALTER TABLE PARTITION BY statements.

For heterogeneous partitioning the first level partitioning type must also be range (single level heterogeneous partitioning is also possible). Dynamic ranges are not supported. Heterogeneous partitioning supports the following additional features:

- Properties (metadata) for partitions; these are defined simply as a value with a boolean state (see *Properties* below)
- Placement of first-level partitions so that partitions can be placed (or moved) directly on a specified host supporting HANA scale-out or Dynamic Tiering solution (placement is also supported for normal balanced partitioning).

The following set of range-range examples illustrate this feature.

Example 1: Heterogeneous 2-level range-range partitions. The SUBPARTITION BY keyword is required to specify the second-level heterogeneous partitions. The syntax in this example shows partitions one and two bracketed together; this is possible if the same subpartition definition applies to both first level partitions.


```
CREATE COLUMN TABLE T_Unbalanced ( BELNR INT, COLGRP NVARCHAR(20) ) PARTITION BY RANGE (COLGRP) (
  ( PARTITION VALUE = 'apple' USING DEFAULT STORAGE,
    PARTITION VALUE = 'google' USING EXTENDED STORAGE )
  SUBPARTITION BY RANGE (BELNR) ( PARTITION VALUE=1, PARTITION VALUE=2 ) )
```

Table 'T_Unbalanced'

#	BELNR	COLGRP
1	1	apple
2	2	google
3	2	google
4	2	apple
5	1	google
6	1	apple
7	2	google
8	2	google
9	1	apple
10	1	google



An additional option which is available with the create table statement is the PRIMARY KEY UPDATE property, this determines if UPDATE statements are allowed on primary key columns. SQL statements are also available to drop heterogeneous subpartitions, redefine heterogeneous first-level partitions, and move heterogeneous first-level partitions (and their subpartitions) to a new location. Full details of these options are given in the *SAP HANA SQL and System Views Reference*.

Properties

The following properties are supported for partitions, values for these are saved in the TABLE_PARTITIONS system view:

Property	Value	Scope
Persistent memory	True / False	Balanced partitioning only
Insert	ON / OFF	Heterogeneous partitioning only
Storage type	default / extended (for multi-store tables)	Heterogeneous partitioning only

These properties can be modified for a partition using the ALTER PARTITION construction:

```
ALTER TABLE <table_ref> ALTER PARTITION <partition_ref> INSERT <ON / OFF>
```

The storage type property only applies at the first level location and can only be modified by using the MOVE PARTITION command:

```
MOVE PARTITION <partition_number> TO <move_location>
```

6.8.9.4 Range Partitioning: More Options

Some special features are available for range partitioning: adding additional ranges, deleting ranges, and using dynamic 'others' partitions.

6.8.9.4.1 Explicit Partition Handling for Range Partitioning

For all partitioning specifications involving range, it is possible to have additional ranges added and removed as necessary. This means that partitions are created and dropped as required by the ranges in use. In the case of multi-level partitioning, the desired operation is applied to all relevant nodes.

Note

If a partition is created and an others partition exists, the rows in the others partition that match the newly-added range are moved to the new partition. If the others partition is large, this operation may take a long time. If an others partition does not exist, this operation is fast as only a new partition is added to the catalog.

Range partitioning requires at least one range to be specified regardless of whether or not there is an others partition. When partitions are dropped, the last partition created cannot be dropped even if an others partition exists.

For range-range partitioning you have to specify whether a partition has to be added or dropped on the first or second level by specifying the partitioning column.

Caution

The DROP PARTITION command deletes data. It does not move data to the others partition.

For more information about the SQL syntax for partitioning, see *SAP HANA SQL and System Views Reference*.

Example

Changing a Table to Add/Drop Partitions

```
ALTER TABLE MY_TABLE ADD PARTITION 100 <= VALUES < 200
ALTER TABLE MY_TABLE DROP PARTITION 100 <= VALUES < 200
```

Example

Changing a Table to Add/Drop an Others Partition

```
ALTER TABLE MY_TABLE ADD PARTITION OTHERS
ALTER TABLE MY_TABLE DROP PARTITION OTHERS
```

You can add a range to the first and second level respectively as follows;

Example

```
CREATE COLUMN TABLE MY_TABLE (a INT, b INT)
```

```
PARTITION BY
  RANGE (a) (PARTITION 1 <= VALUES < 5),
  RANGE (b) (PARTITION 100 <= VALUES < 500)
```

❖ Example

```
ALTER TABLE MY_TABLE ADD PARTITION (a) 5 <= VALUES < 10
ALTER TABLE MY_TABLE ADD PARTITION (b) 500 <= VALUES < 1000
```

❖ Example

```
ALTER TABLE MY_TABLE DROP PARTITION (a) 5 <= VALUES < 10
ALTER TABLE MY_TABLE DROP PARTITION (b) 500 <= VALUES < 1000
```

6.8.9.4.2 Dynamic Range Partitioning

Dynamic Range Partitioning is available to support the automatic maintenance of the others partition.

When you create an others partition there is a risk that over time it could overflow and require further maintenance. Using the dynamic range feature the others partition is monitored by a background job and will be automatically split into an additional range partition when it reaches a predefined size threshold. The background job also checks for empty partitions and if a range partition is found to be empty it is automatically deleted (the others partition is never automatically deleted).

You can create partitions with a dynamic others partition by including the DYNAMIC keyword in the command when you create the partition. This can be either a single level or a second level RANGE partition.

There are three possible ways of setting the row count threshold for dynamic range partitions:

- Using the THRESHOLD keyword in the SQL command (see example below); in this case the value is saved in the table meta data.
- In the DYNAMIC_RANGE_THRESHOLD field of the TABLE_PLACEMENT table which also includes an option to trigger dynamic partitioning.
- A third option is define a threshold value as a system configuration parameter (see *Configuration* below).

❖ Example

This example creates an others partition for table T and then uses the optional keywords DYNAMIC and a <threshold_count> value to specify a maximum size for the partition of 3 million rows. When this threshold is reached a new partition will be created:

```
CREATE COLUMN TABLE T (A VARCHAR(5) NOT NULL, NUM INTEGER NOT NULL)
PARTITION BY RANGE (A AS INT) (PARTITION OTHERS DYNAMIC THRESHOLD 3000000)
```

Not all data types are suitable for dynamic partitioning, the range partitioning column must be a not-nullable column and must be a consistently incrementing numerical sequence such as a timestamp (DATE, TIMESTAMP, SECONDDATE) or an integer (INT, VARCHAR).

You can change the threshold value for an existing dynamic partition or disable dynamic partitioning using the ALTER TABLE command:

❖ Example

The first example here redefines the threshold for the table and the second command turns dynamic range partitioning off:

```
ALTER TABLE T PARTITION OTHERS DYNAMIC THRESHOLD 1000000;
```

```
ALTER TABLE T PARTITION OTHERS NO DYNAMIC;
```

Thresholds and Processing

Tables with a dynamic others partition are monitored by a configurable background job which checks the current size of the partition in comparison to the defined threshold and makes adjustments (either adding a new partition or removing obsolete partitions) as required. Corresponding SQL instructions for these actions are the ADD PARTITION FROM OTHERS (to add a new partition) and DROP EMPTY PARTITIONS (to delete a partition), see `dynamic_range_clause` in the *SAP HANA SQL and System Views Reference* for more details.

The background job runs at a predefined interval and searches for a threshold value in the sequence of options given above starting with the table meta data. Similarly an ALTER TABLE instruction for dynamic partitioning where a threshold is not specified in the SQL command will apply whichever threshold value it finds first.

Configuration

In the [partitioning] section of `indexserver.ini` file the following configuration parameters are available for dynamic partitioning:

- `dynamic_range_default_threshold` - enter the value you require, the default value is 10,000,000 rows. If no other value has been specified then this parameter value is used.
- `dynamic_range_check_time_interval_sec` - the background job runs at a predefined interval defined as a number of seconds. By default this is 900 but this can be changed here if required. You can deactivate this background job by setting the parameter to -1.

6.8.9.5 Partitioning Operations

How a table is partitioned can be determined on creation or at a later point in time. You can change how a table is partitioned in several ways.

You can change partitioning in the following ways:

- Change a partitioned table into a non-partitioned table by merging all of its partitions
- Partition a non-partitioned table
- Re-partition an already-partitioned table, for example:
 - Change the partitioning specification, for example, from hash to round-robin
 - Change the partitioning columns
 - Increase or decrease the number of partitions

Performing a partitioning operation on a table in the above ways can be costly for the following reasons:

- It takes a long time to run, up to several hours for huge tables.
- It has relatively high memory consumption.
- It writes everything to the log (required for backup and recovery).

The re-partitioning operation is a non-blocking process. Apart from a short period at the start as the partitioning process is prepared it does not require an exclusive lock on the database. This means that you can execute SQL DML commands while repartitioning is running, however, DDL operations (table manipulations such as create or drop) are still blocked. At the end of the partitioning process an exclusive lock is again required to apply repartitioning to delta contents.

→ Recommendation

(Re-)partition tables before inserting mass data or while they are still small. If a table is not partitioned and its size reaches configurable absolute thresholds, or if a table grows by a certain percentage per day, the system issues an alert.

→ Recommendation

Although it is possible to (re-)partition tables and merge partitions manually, in some situations it may be more effective to use the redistribution operation available for optimizing table partitioning (for example, if a change of partition specification is **not** required). Redistribution operations use complex algorithms to evaluate the current distribution and determine a better distribution depending on the situation.

i Note

When you change the partitioning of tables the table creation time of the affected tables will be updated to the time you performed the action.

Related Information

[Optimize Table Partitioning \[page 611\]](#)

6.8.9.5.1 Partition a Non-Partitioned Table

You can partition an existing non-partitioned column-store table in the Table Distribution editor of the SAP HANA studio.

Prerequisites

- You have system privilege CATALOG READ and object privileges SELECT and UPDATE either for the table being modified or the schema it is in. If you are the owner of the table, then you can also partition the table without object privileges.

- You have considered the performance impact of partitioning the table.

Procedure

1. Open the *Table Distribution* editor by right-clicking any of the following entries in the *Systems* view and then choosing *Show Table Distribution*:
 - Catalog
 - Schema
 - Tables

A list of all tables is displayed.

i Note

For performance reasons, not all tables are displayed, but only the first 1,000. You can change this setting in the preferences of the SAP HANA studio under **▮ SAP HANA > Runtime > Catalog ▮**. If more tables exist in the selected schema, a message is displayed.

2. Right-click the table that you want to partition and choose *Partition Table*.
3. Specify the first-level partitioning specification:
 - Hash
 - Round robin
 - Range

i Note

Round robin is only available as an option if the table does not have a primary key.

4. Optional: If you want a second level of partitioning, choose *Additional level of partitioning* followed by the required partitioning specification.

The available options depend on the selected first-level specification. The following combinations are possible:

- Hash-range
 - Round-robin-range
 - Hash-hash
5. Enter the information required for the selected partitioning specification(s):

Option	Description
Hash	<ol style="list-style-type: none"> 1. Specify the number of partitions. If the system is distributed, you can enter the number of partitions or choose to have the same number of partitions as the number of hosts. 2. Select the partitioning column(s). For first-level hash partitioning, only primary key columns are available for selection. For second-level hash partitioning, all columns are available. 3. Optional: For date or timestamp columns, specify a date function to partition by year or month.

Option	Description
Round robin	Specify the number of partitions. If the system is distributed, you can enter the number of partitions or choose to have the same number of partitions as the number of hosts.
Range	<ol style="list-style-type: none"> Select the partitioning column. For first-level range partitioning, only primary key columns that have the data type string, integer, and date are available for selection. For second-level range partitioning, all columns with string, integer, and date data types are available. Optional: For date or timestamp columns, specify a date function to partition by year or month. Add the required partitions. For each partition, you must specify the value range, that is the start value and the end value. You can also add partitions for single values. An 'others' partition is created automatically. Choose <i>Validate Input</i> to ensure that the values you have entered are consistent.

i Note

The following general restrictions apply when using the hash and range partitioning specifications:

- The maximum number of partitions supported is 1,000.
- Partitioning columns specified for hash and range must not contain commas (,), dollar signs (\$), or round opening braces (()).
- Ranges must not contain commas (,), semi-colons (;), minus signs (-), asterisks (*), and the pipe character (|).

- Choose *Finish*.

Results

The system partitions the table as specified. The progress of the operation is displayed in the *Progress* view. When partitioning has completed, the information in the *Partition Details for <schema.table>* area is updated accordingly.

i Note

The partitioning operation may take a long time depending on the size of table, available system resources, and so on.

Next Steps

In a multiple-host system, you can now distribute the table partitions to the available hosts. Although you can do this manually, it is recommended that you execute the table redistribution operation *Optimize Table Distribution*.

Related Information

[Optimize Table Distribution \[page 610\]](#)

[Modify Table Distribution Manually \[page 612\]](#)

6.8.9.5.2 Change a Partitioned Table into a Non-Partitioned Table

You can change a partitioned table into a non-partitioned table by merging all of its partitions. You do this in the Table Distribution editor of the SAP HANA studio.

Prerequisites

- You have system privilege CATALOG READ and the object privilege ALTER for the table being modified or the schema it is in.
- If the merge process involves moving some partitions to a single host, the target host must have sufficient memory.
- You have considered the performance impact of merging the table.

Procedure

1. Open the *Table Distribution* editor by right-clicking one of the following entries in the *Systems* view and choosing *Show Table Distribution*:
 - Catalog
 - Schema
 - Tables

i Note

For performance reasons, not all tables of the selected schema are displayed, but only the first 1,000 tables. You can change this setting in the preferences of the SAP HANA studio. If more tables exist in the selected schema, a message is displayed.

2. Right-click the partitioned table that you want to convert to a non-partitioned table and choose *Merge Partitions*.

i Note

When you choose this action the table creation time will be updated to the time you performed the action.

Results

In a single-host system, the system starts to merge the partitions into a non-partitioned table immediately.

In a multiple-host system, the system first checks that all partitions reside on the same host. If this is not the case, you are prompted to select the host to which you want to move them all. Before moving the partitions to the selected host, the system checks that the host has sufficient memory. If this is the case, the system first moves the partitions before merging them.

The progress of the operation is displayed in the *Progress* view. When the partitions have been completely merged, or moved and merged, the information in the *Partition Details for <schema.table>* area is updated accordingly.

i Note

Merge operations may take a long time depending on the size of the partitions, whether or not the partitions have to be moved first, available system resources, and so on.

6.8.9.6 Time Selection Partitioning (Aging)

The SAP HANA database offers a special time selection partitioning scheme, also called aging. Time selection or aging allows SAP Business Suite application data to be horizontally partitioned into different temperatures like hot and cold.

SAP Business Suite ABAP applications can use aging, which must not be used for customer or partner applications, to separate hot (current) data from cold (old) data by using time selection partitioning to:

- Create partitions and re-partition
- Add partitions
- Allocate rows to partitions
- Set the scope of Data Manipulation Language (DML) and Data Query Language (DQL) statements.

Setting the DML and DQL scope is the most important aspect of time selection partitioning. It uses a date to control how many partitions are considered during SELECT, CALL, UPDATE, UPSERT and DELETE. This date may be provided by the application with a syntax clause and it restricts the number of partitions that are considered.

For example a SELECT statement may be issued that retrieves all data having a date greater or equal to May 1st, 2009. It shall also include the current/hot partition. On the other hand, UPDATE operations can also be restricted in the same way. If a date is provided, the current partition is also always included.

Unique Constraints for Cold Partitions

By default SAP HANA enforces unique constraints on all partitions. The application may actively overrule this behavior for cold partitions though. This requires that the applications enforce uniqueness for cold partitions by themselves. Duplicate keys are then considered to be application errors.

The reason is that typical OLTP workload in SAP Business Suite for SAP HANA is executed on the current/hot partition and its performance shall not be affected by unique checks for cold partitions that are not relevant for typical OLTP processing.

If the application overrules the uniqueness constraints:

- A row of the current/hot partition may be in conflict with a row of a cold partition,
- A row of a cold partition may be in conflict with a row of another cold partition, and
- A row within a cold partition may be in conflict with another row within the same cold partition.

Partitioning is transparent from an SQL perspective. If a table has a unique index or a primary key and if it has duplicates in cold partitions, a SELECT may return duplicates for a unique index or primary key. This behavior is correct from a database perspective, but this is considered an application error. The database will return an undefined result set. The only kind of statement that will return a correct result set if duplicate primary keys exist is a SELECT statement, which does nothing but select data with a WHERE clause on the full key (no joins, aggregations, aggregate functions or the like and not complex WHERE conditions). There is no guarantee with respect to the result set for further unique constraints if duplicates exist.

Paged Attributes

Cold partitions may optionally be created as paged attributes. This reduces memory consumption. Memory for resident pages are included in the system views M_CS_TABLES, M_CS_COLUMNS and M_CS_ALL_COLUMNS. The field MEMORY_SIZE_IN_MAIN and related statistics include both the paged and non-paged memory for tables or columns.

Global statistics for resident pages and can be found in the M_MEMORY_OBJECT_DISPOSITIONS view. The number and size of pages used by paged attributes are tracked in PAGE_LOADABLE_COLUMNS_OBJECT_COUNT and PAGE_LOADABLE_COLUMNS_OBJECT_SIZE, respectively.

Converting Aging tables to Column Store Tables

In some specific cases it is possible to convert time selection partitioned tables to column store tables.

- Single-level aging partitioned table can be converted to non-partitioned table
- Two-level aging partitioned table can be converted to one-level partitioned table with second-level time selection partitions merged
- Single-level or two-level aging multistore tables can be converted to regular column store partitioned tables by firstly moving partitions from extended storage to the column store and then converting. Refer to the topic *Convert a Multistore Table to a Partitioned Column Store Table* in the SAP HANA Dynamic Tiering documentation set for more details.

Examples to illustrate this are given here showing firstly the command to create the table and then alter it:

❁ Example

1) Convert single-level aging table and merge the partitions to an unpartitioned column store table

```
CREATE COLUMN TABLE TAB (A INT, B INT PRIMARY KEY, _DATAAGING NVARCHAR(8))
PARTITION BY RANGE (_DATAAGING)
(USING DEFAULT STORAGE (PARTITION value = '00000000' IS CURRENT,
PARTITION '20100101' <= VALUES < '20110101', PARTITION OTHERS))
WITH PARTITIONING ON ANY COLUMNS ON
FOR NON CURRENT PARTITIONS UNIQUE CONSTRAINTS OFF
FOR DEFAULT STORAGE NON CURRENT PARTITIONS PAGE LOADABLE;
```

```
ALTER TABLE TAB MERGE PARTITIONS;
```

2) Convert two-level aging-partitioned table to a single-level column store partitioned table (HASH-RANGE to HASH)

```
CREATE COLUMN TABLE TAB (A INT, B INT PRIMARY KEY, _DATAAGING NVARCHAR(8))
PARTITION BY HASH(B) PARTITIONS 2, RANGE (_DATAAGING)
(USING DEFAULT STORAGE (PARTITION value = '00000000' IS CURRENT,
PARTITION '20100101' <= VALUES < '20110101', PARTITION OTHERS))
WITH PARTITIONING ON ANY COLUMNS ON
FOR NON CURRENT PARTITIONS UNIQUE CONSTRAINTS OFF
FOR DEFAULT STORAGE NON CURRENT PARTITIONS PAGE LOADABLE;
```

```
ALTER TABLE TAB PARTITION BY HASH (B) PARTITIONS 2;
```

Note that in this case the partition number must be 2.

3) Convert two-level aging-partitioned table to single-level column store partitioned table (RANGE-RANGE to RANGE)

```
CREATE COLUMN TABLE TAB (A INT PRIMARY KEY, B INT, _DATAAGING NVARCHAR(8))
PARTITION BY RANGE(A) (PARTITION VALUE = 1, PARTITION 10 <= VALUES < 20,
PARTITION OTHERS), RANGE (_DATAAGING)
(USING DEFAULT STORAGE (PARTITION value = '00000000' IS CURRENT,
PARTITION '20100101' <= VALUES < '20110101', PARTITION OTHERS))
WITH PARTITIONING ON ANY COLUMNS ON
FOR NON CURRENT PARTITIONS UNIQUE CONSTRAINTS OFF
FOR DEFAULT STORAGE NON CURRENT PARTITIONS PAGE LOADABLE;
```

```
ALTER TABLE TAB PARTITION BY RANGE(A) (PARTITION VALUE = 1, PARTITION 10 <=
VALUES < 20, PARTITION OTHERS);
```

Note that in this case the partition spec must be the same as the first-level spec of the aging table.

More information on data aging is available in the following SAP Note: 2416490 - FAQ: SAP HANA Data Aging in SAP S/4HANA.

Related Information

[SAP Note 2416490](#)

6.8.9.7 Partitioning Consistency Check

You can call general and data consistency checks for partitioned tables to check, for example, that the partition specification, metadata and topology are correct.

There are two types of consistency checks available for partitioned tables:

1. General check
Checks the consistency among partition specification, metadata and topology.
2. Data check
Performs the general check and additionally checks whether all rows are located in the correct parts.

To perform the general check, execute the following statement:

```
CALL CHECK_TABLE_CONSISTENCY('CHECK_PARTITIONING', '<schema>', '<table>')
```

To perform the extended data check, execute:

```
CALL CHECK_TABLE_CONSISTENCY('CHECK_PARTITIONING_DATA', '<schema>', '<table>')
```

If any of the tests encounter an issue with a table, the statement returns a row with details on the error. If the result set is empty (no rows returned), no issues were detected.

If the extended data check detects, that rows are located in incorrect parts, this may be repaired by executing:

```
CALL CHECK_TABLE_CONSISTENCY('REPAIR_PARTITIONING_DATA', '<schema>', '<table>')
```

i Note

The data checks can take a long time to run depending on the data volume.

Related Information

[Table Consistency Check \[page 509\]](#)

6.8.9.8 Designing Partitions

There are a number of factors to consider to optimize the design of your data partitioning strategy including how it will affect select and insert performance and how it will adjust to data changes over time.

Different partitioning strategies need to be tested to determine the best one for your particular scenario. Based on your tests you should choose the partitioning strategy that shows the best performance for your scenario. The design principals listed here are aids to help you decide on the correct partitioning strategy for your scenario.

i Note

SAP Business Warehouse on SAP HANA handles partitioning itself. Do not interfere with its partition management unless this has been recommended by SAP.

Query Performance

- For replicated dimension tables the database tries to use replicas that are local to the fact table partitions.
- Partition pruning analyzes the WHERE clauses and seeks to reduce the number of partitions. Try to use partitioning columns that are often used in WHERE clauses. This reduces run time and load.
- Usually hash partitioning is the best partitioning scheme for the first level, especially in scale out scenarios. This is because the client may already use pruning on the client machine and send the query directly to the host that holds the data, where possible. This is called “client-side statement routing”. This is especially important for single select statements.
- Use as many columns in the hash partitioning as required for good load balancing, but try to use only those columns that are typically used in requests. In the worst case only single select statements may leverage pruning.
- If tables are joined with each other, it is beneficial if the tables are partitioned over the same columns and have the same number of partitions. This way the join may be executed locally in scale out scenarios and the network overhead is reduced.
 - This guarantees that the matching values are in a partition with the same part ID. You have to put all parts with the same ID on the same host.
- Queries do not necessarily become faster when smaller partitions are searched. Often queries make use of indexes and the table or partition size is not significant. If the search criterion is not selective though, partition size does matter.

Data Manipulation Language (DML) Performance

- If insert performance is key to your scenario, a larger number of partitions might show better results. On the other hand, a higher number of partitions may reduce query performance.
- Partition pruning is used during DML operations.
- For replicated column store tables, all DML operations are routed through the host with the master partition (where the replica with Part ID 1 is located).
- If there is a unique constraint on a non-key column, the performance will suffer exponentially with the number of partitions on other servers. This is because the uniqueness on all partitions has to be checked. Therefore, if partitioning is required, consider a low number of partitions and ideally put all partitions on the same host. This way the number of remote calls is reduced.

Data Lifecycle

If time-based partitioning is suitable for the dataset being partitioned, it should always be used as it has a number of advantages:

- The runtime of a delta merge is dependent on the size of the main index. This concept leverages the fact that new data is inserted into new partitions whereas data in old partitions is infrequently updated. Over time, the formerly new partitions become old and new partitions are being created. Delta merges on old partitions are not required anymore. This way the overall runtime of delta merges does not increase with the table size, but remains at a constant level. Using time-based partitioning often involves the use of hash-range partitioning with range on a date column. This requires knowledge of the actual values for range partitioning.

- By using explicit partition management, new partitions can be created, for example, one partition per calendar week and old partitions may be dropped entirely rather than deleting individual rows.
- If you split an index, always use a multiple of the source parts (for example 2 to 4 partitions). This way the split will be executed in parallel mode and also does not require parts to be moved to a single server first.
- Do not split/merge a table unless necessary. These operations write all data into the log which consumes a high amount of disk space. Moreover, the operations take a long time and locks the table exclusively (only selects are allowed during partitioning operations). ADD PARTITION can be used to add additional partitions. If there is no 'others' partition, this call only creates a new partition which is fast and happens in real time after an exclusive lock of the table was acquired. On the other hand, if the table has an others partition, a call to ADD PARTITION causes the existing others partition to be split into a new others partition and newly requested range. This is a costly operation. Therefore it is recommended, that if ADD PARTITION is used frequently in a scenario, the table shall not have an others partition.

Partition Size

- Due to the high number of factors to be considered when evaluating a partitioning scheme a recommendation for partition sizes cannot be provided. If you do not know if you will partition a table at all or with how many partitions you need to start with measurements. Here are some suggested starting points:
 - If a table has less than 500 million rows, do not partition it at all unless:
 - The corresponding tables in joins are partitioned. If they are try to find mutual partitioning columns.
 - Table growth is expected. Since re-partitioning is time consuming, it is recommended to split a table while it is still small.
 - If your table has more than 500 million rows, choose 300 million per partition.
 - Keep in mind that a partition must not exceed 2 billion rows.
- Be aware that a higher number of partitions might lead to higher memory consumption as each partition has its own exclusive dictionary, which is not shared. If each partition stores disjunctive values, this is not an issue. On the other hand, if each partition has similar or the same values this means that the dictionaries have similar data which is stored redundantly. In this case consider using fewer partitions

Table Design

- If the data is replicated into the SAP HANA database, it might be fair from a data consistency perspective to remove a primary key or to extend the key (since the key constraint is enforced in the source database). This way you might be able to have multiple tables with the same partitioning columns even though the original database design would not have allowed it. Having the same partitioning columns is ideal as related data may reside on the same physical host and therefore join operations may be executed locally with no or hardly any communication costs.
- When designing database schemas for dependent hosts, for example, a database structure for business objects with header and leaf nodes, do not use a single GUID column as the primary key. In such a case it is hardly possible to have all related data (for example, a business object instance) on the same host. One option might be to have a GUID as the primary key in the header table and each host, irrespective of its level, could have that GUID as the first primary key column.

- Do not define a unique constraint on a partitioned table unless absolutely necessary.
- On the second partitioning level, a non-primary key column may be used. Still, the unique constraint has to be enforced on all parts of the respective first-level partition. Since all parts of one first-level partition are moved as a whole, this unique check is always local.
- In case the database table is replicated from another database, an others partition for range is generally recommended. If a proper range is not defined, the insert statement will fail and the data will not get replicated properly.
- Ideally tables have a time criterion in the primary key. This can then be used for time-based partitioning. Number ranges and so on can also be used. The advantage of number ranges is that it is easy to form equally sized partitions, but on the other hand it introduces an administrative burden the amount of data that is loaded needs to be closely monitored and new partitions need to be created in advance. In case of actual dates, you only need to periodically create new partitions, for example, before a new quarter starts.

Other Considerations

- Use `GET_NUM_SERVERS()` in scripts for hash and round-robin partition specifications. This way Table Placement is used to calculate the number of partitions that will be used in the given landscape.
- If it is likely that a table has to be re-split in future and range partitioning is used, define an others partition. (If it is not defined upon table creation, it can be created afterward and if required dropped after the split operation).
- To check whether a table is partitioned, do not consider the existence of a partition specification in the metadata. Instead check `IS_PARTITIONED` in `M_TABLES` or for the existence of parts, for example in `M_CS_TABLES`. It is allowed that a partition specification is defined which does not immediately lead to a partitioned table.

6.8.9.8.1 Static and Dynamic Partition Pruning

An important partitioning strategy to improve performance is to match partitions wherever possible with the most frequently queried data so that data pruning is possible.

Pruning takes place automatically in the background and tries to eliminate any unnecessary partition from the selection which is not essential for the query result. If your partitions are designed to support this, pruning can avoid accessing and loading into memory partitions which are not required which reduces the overall load on the system. The classic use case for effective pruning is where date-based partitions are used. For example, if a table is partitioned by year, a query restricted to the data of a single year is executed only on the partition with data for the selected year.

Two forms of partition pruning are possible: static and dynamic.

- **Static partition pruning** is based on the partition definition. The query optimizer analyzes the WHERE clause of queries to determine whether or not the filters match the given partitioning specification of a table. If a match is found, it may be possible to target only the specific partitions that hold the data and thus reduce the number of partitions being queried.
- **Dynamic partition pruning** is content-based and can be applied to historical partitions of a HANA aging table. This type of pruning takes place at run time but the decision about which partitions to load is based on pre-calculated column statistics of the data in selected columns. Based on the existence of statistics

valid for partition pruning (which must be up to date at the time the query is processed), query evaluation can determine if it is necessary to read the partitioned column of data and load it into memory.

Statistics for Dynamic Partition Pruning

Statistics for dynamic partition pruning are created and maintained as part of the SAP HANA data statistics functionality and are maintained and used only in the indexserver instance where the partition resides. Statistics are explicitly created and refreshed through SQL statements (such as CREATE STATISTICS) and to use dynamic pruning you must first run CREATE STATISTICS for a specified table, partition and column (see following *Enabling* subsection). Any column other than the partition key can be used. The statistics for dynamic partition pruning are persisted with other statistic values in the system catalogs (this takes place asynchronously in the background), however, the pruning statistics values are also cached in memory for the pruning process.

Statistics for pruning have a limited life-span: as new data is added to the column the content of these statistics becomes stale. It is therefore necessary to periodically refresh the pruning statistics. This is done automatically each time delta merge runs and can also be done on demand (see DDL statements below). The statistics are considered to be invalid as soon as new data is added to the column and pruning is then no longer applied to the column until the pruning statistics have been refreshed.

With historically aged data partitions, it is likely that the greatest number of inserts and deletions occurs in the current time period. The older time partitions will be more stable, and for this reason dynamic pruning is only applied to time-based selections on historical data not the current time period.

Enabling Dynamic Partition Pruning

Static and dynamic pruning are designed as integral features of the query optimizer and execution processes. However, not all query engines are currently able to take full advantage of dynamic pruning and the setting **use_dynamic_pruning** in the query_mediator section of the indexserver.ini file must be set to **True** to enable dynamic pruning for the HANA OLAP engine plan.

Statistics for dynamic partition pruning must be explicitly created for a specific table using the CREATE STATISTICS command as shown in the following example (also applies to ALTER STATISTICS). This creates statistics for all partitions:

```
CREATE STATISTICS test_dynPruning_tabl_coll
    ON tabl (coll)
    TYPE SIMPLE VALID FOR DATA DEPENDENCY;
```

For dynamic partition pruning the statistics object type 'SIMPLE' is required with the property set 'VALID FOR DATA DEPENDENCY' as shown here. It is only applicable to partitioned column store tables and multistore tables using time selection partitioning. Only the following datatypes are supported for the statistics columns: VARCHAR (strings with numerical content), INTEGER, DECIMAL, DATE.

Executing the CREATE STATISTICS statement in this way initializes the column for dynamic pruning and populates the pruning statistics. The pruning statistics values are recalculated either by delta merge or, additionally, an explicit refresh statement is available to recalculate the statistics on demand:

```
REFRESH STATISTICS test_dynPruning_tabl_coll
```

Similarly, a DROP STATISTICS command is available and can be used in the same way to drop pruning statistics for a table.

Refer to the *SAP HANA SQL and System Views Reference* for full details.

Monitoring Views and Tracing

Information about statistics for dynamic partition pruning is available in the M_DATA_STATISTICS view.

Diagnostics information is also available. You can see debug details of the pruning process by setting the trace level of the `part_pruning` configuration parameter (in the trace section of the `indexserver.ini` file) to `debug`.

See also *Configure Traces* for more information.

Related Information

[Configure Traces in SAP HANA Studio \[page 683\]](#)

[Time Selection Partitioning \(Aging\) \[page 561\]](#)

6.8.9.8.2 Creating an Effective Partitioning Scheme

This checklist demonstrates how to choose a good partitioning scheme for given tables.

- Tables of above 500 million rows are good candidates for partitioning. This also applies to small tables that are often joined with tables of above 500 million rows.
- If the table has a unique index (other than the primary key), the table may be partitioned, but the additional unique checks introduce a performance penalty.
- Check the primary key.
 - If none exists, any columns may be used for Hash partitioning.
 - If one is present, identify the minimal set of columns that are required to have equally balanced partitions; a sufficiently high number of distinct values is required. Keep in mind that if these columns are all in the WHERE clause of a query, partition pruning may be leveraged.
 - In the case of tables that are replicated into SAP HANA, it may be legitimate to drop the primary key since it is checked in the source database.
- Take other tables into consideration that are often used in joins with the current table. Ideally they have the same number of partitions and partitioning columns.
- Identify time-based attributes; this may be a date, year or at least a sequence. Use them for time-based partitioning. Ideally this column is part of the primary key.
- If you define range partitioning, decide whether or not you require an others partition. Ideally, no others partition is required.
- Decide on the number of partitions. Use Table Placement rules, if applicable.
- In case of a scale out system, move all corresponding partitions to the respective hosts.
- Run extensive performance tests with the most-prominent queries and/or DML load. Try to use analytical views. Vary partitioning columns, partitioning schemes and the number of partitions.

6.8.9.8.2.1 Partitioning Example

This example describes an initial and subsequently improved partitioning schema for a database storing test results.

Assume that for each make various tests run. So many rows have to be stored for a TEST_CASE and MAKE_ID.

Original Table Design

- The original table design suggested having a sequence number as the only primary key column. There are time columns marking the start and end of the test run.
- The table was partitioned by Hash over the sequence number and by range over the start date.

Access Pattern

There are two prominent ways how the data is accessed:

1. Select all test results for a make (“Did my make pass the tests?”)
2. Select results of a single test for all makes within the last month (“How is the history of my test? Has an error happened before?”)

So typically either the TEST_CASE or the MAKE_ID is in the WHERE clause, sometimes both when investigating details of a test run.

Problems with this Original Table Design and Partitioning

- The sequence number is not often part of the WHERE clause and hence all partitions are considered in the query. This is especially an issue in scale-out landscapes where OLTP-like queries are ideally only executed on a single node.
- There is a hash-range partitioning with range on a date column. This allows time-based partitioning to be used. But the date column is not part of the primary key. Therefore the unique constraint on the primary key has to be ensured by checks with the parts of the first-level partition.

Suggested Table Design

- Have TEST_CASE, MAKE_ID and SEQ_NUMBER as the primary key. The actual requirement that there are uniquely identifiable rows for a combination of TEST_CASE and MAKE_ID is met.
- Partition by hash-range with hash over TEST_CASE and range over MAKE_ID. The MAKE_ID increases over time and therefore is also a good column to use for time-based partitioning.

Reasoning

- No primary key checks with other partitions (using range on primary key column).
- Good pruning since partitioning columns match the typical access pattern.
 - If the query has the MAKE_ID in the WHERE clause (compare query type 1), all servers that hold partitions have to be considered, but only a single partition per server.
 - If the query has the TEST_CASE in the WHERE clause (compare type 2), only one server has to be considered (compare Client-Side Statement Routing), but all partitions on that server.
 - If MAKE_ID and TEST_CASE are in the WHERE clause, only a single partition on one server has to be considered.

In this scenario there is one type of query which will cause that all servers that hold partitions are considered. This is not ideal, but cannot always be prevented depending on the nature of the data and access patterns.

6.8.9.8.3 Partitioning Limits

General restrictions that apply to the use of partitioning are explained here.

General Limitations

- The maximum number of partitions for one table is 16000. A table may be re-partitioned as often as required. The limit of 16000 partitions is independent from the location of the partitions in a distributed (scale-out) landscape.
- Column names containing spaces are not supported for partitioning.
- Partitioning columns specified for hash and range partitioning must not contain commas (","), dollar signs ("\$\$") or round opening parentheses ("(").
- Range partition values must not contain commas (","), minus signs ("-"), asterisks ("*") or the space character.
- When using an equidistant series and table partitioning, for efficient compression ROUND ROBIN partitioning should not be used. HASH or RANGE partitioning should be used so that records with the same series key are in the same partition.

History Tables

Tables with history tables can also be partitioned. A history table is always partitioned with the same partitioning type as the main table.

- The 2 billion rows barrier is also valid for the history tables and in case of partitioned history tables this will be also hold true on a per partition basis.
- If a table uses multi-level partitioning, it is possible to use partitioning columns which are not part of the primary key. This feature cannot be used in conjunction with history tables.

- Tables with history cannot be replicated.

Data Types

The following restrictions apply to specific data types:

- Hash partitioning: Only the following column store data types are allowed in the partitioning columns: TINYINT, SMALLINT, INT, BIGINT, DECIMAL, DECIMAL(p,s), CLOB, NCLOB, SHORTTEXT, VARCHAR, NVARCHAR, BLOB, VARBINARY, DATE, TIME, TIMESTAMP and SECONDDATE. Memory LOBs (ST_MEMORY_LOB type only) are supported but not disk LOBs.
- Range partitioning: Only the following column store data types are allowed in the partitioning column: STRING, TINYINT, SMALLINT, INT, SHORTTEXT, VARCHAR, NVARCHAR, DATE, TIME, TIMESTAMP, SECONDDATE, FIXED, RAW (SQL Binary/Varbinary). Memory LOBs (ST_MEMORY_LOB type) are supported but not disk LOBs.

i Note

The time unit used by partitioning is the whole day, it is not possible to consider smaller units such as hours or minutes.

Partitioning Columns

- The data type of columns used as partitioning columns must not be changed.
- If a table has a primary key, a column must not be removed from the primary key if it is being used as partitioning column. In the case of a multi-level partitioning, this applies to the first level. It is always possible to remove the entire primary key.

Related Information

[SAP Note 2044468](#) 

6.8.9.9 System Views for Monitoring Partitions

A number of system views allow you to monitor your partitions.

System View	Description
TABLE_PARTITIONS	Contains partition-specific information for partitioned tables including details of level 1 and level 2 ranges.
PARTITIONED_TABLES	General partitioning information for all partitions of a table

System View	Description
M_TABLE_PARTITIONS	Contains column table information per partition.
TABLE_COLUMNS	Contains table column information.
M_CS_TABLES	Shows run time data per partition (the PART_ID value is the sequential number of the partition). Be aware that after a split/merge operation the memory size is not estimated and therefore the values show zero. A delta merge is required to update the values.
M_TABLES	Shows row counts and memory usage in an aggregated way for partitioned tables (IS_PARTITIONED set to TRUE). Information is based on M_CS_TABLES.
M_CS_PARTITIONS	Shows which partitions or sub-partitions in column store tables form a partition group. This information is for example required when partitions or groups of sub-partitions are to be moved to another host.
M_TABLE_PARTITION_STATISTICS	Shows selection statistics for each partition. In particular, it shows the number of times the partition is selected and includes the timestamp of the last time it was selected. This view is not populated by default and can be enabled through configuration parameter <code>partition_statistics_select_enabled</code> in the <code>partitioning</code> section of the <code>indexserver.ini</code> file.

6.8.9.10 Multistore Tables

A multistore table is a range partitioned SAP HANA column table that has at least one partition in SAP HANA and other partitions in different physical stores associated with SAP HANA.

Multistore capability is currently implemented through SAP HANA dynamic tiering, which provides extended storage as an alternate physical store for SAP HANA table partitions. Multistore tables are managed and monitored just like other columnar tables, but offer the advantages of disk-based storage for data aging through extended storage.

Multistore data management allows flexible administration, such as:

- Moving data between extended or default storage.
- Creating or dropping partitions directly in either extended or default storage.
- Repartitioning a table, if new partitioning does not move data between default and extended storage.

Data modification operations let you easily change the storage type for data. For example, you can change the storage type for a partition when you import data into a table or modify the data in a multistore table during an insert operation.

i Note

Multistore functionality does not support all features of SAP HANA. For details of features supported, see *Multistore Tables* in the *SAP HANA Dynamic Tiering Administration Guide*.

6.8.10 Table Placement

Table classification and table placement configuration, enhanced by partitioning, build the foundation for controlling the data distribution in a SAP HANA scale-out environment.

Table Classification and Placement

Application data is usually stored in a multitude of database tables, and data from several of these tables is combined via SQL operations like join or union when it is queried. As these relations between different tables are defined in the application code, this information is not available within SAP HANA. The table classification feature provides a possibility to push down this semantic information in the database by allowing administrators to define groups of tables. This information can be used, for example, when determining the number of partitions to be created, or, in the case of a scale-out landscape, the node where to locate the tables or partitions.

The classification is performed by providing each table with a group name, group type and subtype. Based on combinations of these elements as well as the table names and schema names, a set of configuration values can be defined as table placement rules, the rules are used to control, for example, the placement of partitions or the number of partitions during operations like table creation or redistribution. By doing this, associated or strongly related tables are placed in such a way that the required cross-node communication is minimized for SQL operations on tables within the group.

Table placement rules are applied during system migration or table creation, but it may also be necessary to adjust the location or the number of partitions on an ongoing basis for handling data growth. Table redistribution can therefore also be run on demand to optimize the landscape as the system evolves. Repartitioning is always necessary, for example, for any table or partition in the database which reaches the maximum count of 2 billion rows.

The following tools are available to perform table repartitioning and redistribution; these tools evaluate the current landscape and determine an optimized distribution:

- SAP HANA Table Redistribution (available in SAP HANA cockpit, SAP HANA studio, or executed from the command line)
- Data Distribution Optimizer (part of SAP HANA Data Warehousing Foundation)

Balancing a SAP HANA scale-out landscape with these tools is done in two stages:

1. Generation of a plan based on table placement rules (described in detail below). After generating the plan, you can review it and adjust the definition of the rules if required.
2. Execution of the plan which implements the partitioning and distribution changes.

Because split table and move table are operations that require table locks, the execution of the plan should not be performed during a period where there is heavy load on the database.

i Note

Please refer to the corresponding documentation for details. In particular, the *SAP HANA SQL and System Views Reference* gives details of all SQL syntax for the operations described here.

Table Classification and Table Placement Rules

Table placement rules are defined in the table SYS.TABLE_PLACEMENT; system privilege TABLE ADMIN is required to maintain these settings. Placement rules basically address the following areas:

- Classification: related tables which must be located together are organized in groups
- Configuration settings to manage partitioning (number of initial partitions, split threshold and so on)
- Physical distribution or location of tables or partitions in the server landscape.

When creating a table, if the defined rules match with a table or a table group, SAP HANA will consider them while creating the table.

Please keep in mind that partition specifications still need to be defined by the application.

Related Information

[Scaling SAP HANA \[page 1404\]](#)

[Data Temperature: Extension Nodes \[page 581\]](#)

[Redistributing Tables in a Scaleout SAP HANA System \[page 600\]](#)

6.8.10.1 Table Classification (Groups)

Associated tables can be classified by a common table group.

The SQL interface of SAP HANA provides three possible kinds of classifications: group name, group type and subtype. Tables which have been classified with group information are included in the SYS.TABLE_GROUPS table and you can review classification details in the monitoring view SYS.TABLE_GROUPS (details below). Tables with the same group name are kept on the same host, or, in the case of partitioned tables that are distributed over several hosts, corresponding first-level partitions are distributed for all tables the same way.

One table in the group is defined as the leading table and its table placement settings are applied to all other tables in the group. This can be, for example, the location, or in the case of partitioned tables (if SAME_PARTITION_COUNT is set in SYS.TABLE_PLACEMENT, see below), the number of first-level partitions.

Note that specific applications, like SAP BW, classify objects automatically as they are created and must not be changed manually.

SYS.TABLE_GROUPS

Column	Description
SCHEMA_NAME	The schema name.
TABLE_NAME	The table name.
GROUP_NAME	The group name
GROUP_TYPE	The group type. Example: in SAP BW there are predefined group types that classify the tables associated with a BW object, such as: sap.bw.cube (InfoCubes), sap.bw.dso (DataStore Objects), sap.bw.psa (PSA tables) and so on.

Column	Description
SUBTYPE	The subtype. This is required for some group types, for example in SAP BW: <ul style="list-style-type: none"> a table belonging to an InfoCube (group type sap.bw.cube) can be a fact table (subtype FACT_IMO) or a dimension table (subtype DIM) a table belonging to a DataStore Object (group type sap.bw.dso) can be an active table (subtype ACTIVE), an activation queue (subtype QUEUE), or a Changelog (subtype CHANGE_LOG)
IS_GROUP_LEAD	Determines the leading table within a group. If none is set, the largest, partitioned, non-replicated column store table is used as leading table.

❖ Example

In this example the following three tables that belong to SAP Business Warehouse (BW) DataStore Object with technical name ZFIGL, all have the same group name value and form the group 'ZFIGL'.

Example Table Group

Table Name	Group Type (GROUP_TYPE)	Subtype (SUBTYPE)	Group Name (GROUP_NAME)
/BIC/AZFIGL00	sap.bw.dso	ACTIVE	ZFIGL
/BIC/AZFIGL40	sap.bw.dso	QUEUE	ZFIGL
/BIC/B0000197000	sap.bw.dso	CHANGE_LOG	ZFIGL

For native applications, the application developer can define a grouping manually, for example, by grouping tables together that are often joined. The following statements show examples of setting table group attributes using CREATE and ALTER table:

```
CREATE COLUMN TABLE "TEST".A1 (A INT) GROUP TYPE ABC GROUP SUBTYPE T
```

```
ALTER TABLE "SAPLG1"."/BIC/MUCSTR000000" SET GROUP NAME "ABC"
```

Creating Groups Dynamically

There are a number of options for creating table groups dynamically based on the information available in the SQL Plan Cache.

- The Join Path Analysis tool is available within the Data Distribution Optimizer, or the ABAP grouping report for S/4HANA scale-out.
- Table Redistribution also includes an optional preparation step to integrate the Group Advisor tool into the plan generation process to create table groups dynamically.

Related Information

[Configuration of Table Redistribution \[page 618\]](#)

6.8.10.2 Table Placement Rules

The TABLE_PLACEMENT table provides a customizing interface which can be used for the dynamic management of partitions and locations.

The TABLE_PLACEMENT Table

The structure of the table is shown here. Parameters 6-10 define how a table or a group of tables is partitioned; further information about usage is provided in the subsections which follow.

SYS.TABLE_PLACEMENT

#	Column	Description
1	SCHEMA_NAME	The schema name (as for table groups).
2	TABLE_NAME	The table name.
3	GROUP_NAME	The group name.
4	GROUP_TYPE	The group type.
5	SUBTYPE	The subtype.
6	MIN_ROWS_FOR_PARTITIONING	Partitioning rule: the number of records that must exist in the table before the number of first-level partitions is increased above 1.
7	INITIAL_PARTITIONS	Partitioning rule: determines the number of initial partitions to create, for example, HASH 1, HASH 3.
8	REPARTITIONING_THRESHOLD	Partitioning rule: if the row count exceeds this value then further split iterations are considered.
9	DYNAMIC_RANGE_THRESHOLD	Applies to tables that use the dynamic range partitioning feature. Overwrites the system default value defined in indexserver.ini > [partitioning] > dynamic_range_default_threshold (10,000,000) for that specific combination of schema / table / group characteristics.
10	SAME_PARTITION_COUNT	Specifies that all partitions of the tables in a group will contain the same number of partitions. Globally maintained in global.ini > [table_placement] > same_num_partitions but in case of several applications with deviating settings, it can be maintained on a more granular level.
11	LOCATION	Location rule: master, slave, all (see below).

Partitioning Rules

The partitioning parameters are used to define how a table or a group of tables is partitioned if the table has a first-level partitioning specification of HASH or ROUNDROBIN. RANGE partitioning is not handled in this way.

If the number of rows is lower than `MIN_ROWS_FOR_PARTITIONING`, the table consists only of one partition. If this minimum row limit is exceeded, the table will be partitioned in as many parts as to fulfill the following constraints:

- The number of partitions is larger or equal to the value of $(\text{row count of the table}) / \text{REPARTITIONING_THRESHOLD}$
- The number of partitions is a multiple of `INITIAL_PARTITIONS`
- The number of partitions is smaller or equal to the number of hosts in case parameter `max_partitions_limited_by_locations` is not set to false and the number of partitions is less than the value of parameter `max_partitions` (see below).

Therefore, if the table has more than one partition, there are at least `INITIAL_PARTITIONS` partitions, and each partition has less than `REPARTITIONING_THRESHOLD` records. Here partitions refer to first-level partitions (of type HASH or ROUNDROBIN).

Please note, when a partitioned table is created without an estimated row count (default behavior) a partitioned table is created with `INITIAL_PARTITIONS` first-level partitions, whereas in a redistribution it is targeted to have a single first-level partition (assuming `MIN_ROWS_FOR_PARTITIONING > 0`). In specific applications, creation is performed with an estimated row count, for example, BW with 1 million and therefore it will be created with only one first level-partition (assuming `MIN_ROWS_FOR_PARTITIONING > 1,000,000`).

Repartitioning

There is no automatic repartitioning when threshold values are exceeded. Instead, this is proposed the next time the redistribution process is executed.

The values entered for partitioning must be consistent with the physical landscape, especially the number of server nodes available:

- If repartitioning is necessary, tables are only repartitioned by doubling the number of existing (initial) partitions. This is done for performance reasons. The maximum number of (first-level) partitions reached by that process is defined by parameter `global.ini > [table_placement] > max_partitions` (default 12).
- By default, the system does not create more partitions than the number of available hosts (or more specifically possible locations). For example, if `INITIAL_PARTITIONS` is set to 3 but the distributed SAP HANA database has five possible locations, repartitioning from three to six partitions would not take place. A table can have more than one partition per host if the parameter `global.ini > [table_placement] > max_partitions_limited_by_locations` is set to false (default true). This rule is disregarded if a higher number of first level partitions is required to partition groups with more than 2 billion records (`global.ini > [table_placement] > max_rows_per_partition`, default = 2,000,000,000).

Location

There are predefined values for possible locations:

- master: represents the master node
- slave (or slaves): represents all slave nodes that belong to the worker group 'default'
- all: represents all nodes that belong to the worker group 'default', i.e. master node and slave nodes.

The worker group assignment can be found in entry `WORKER_ACTUAL_GROUPS` of view `M_LANDSCAPE_HOST_CONFIGURATION` and accessed by executing the following procedure:

```
call SYS.UPDATE_LANDSCAPE_CONFIGURATION( 'GET WORKERGROUPS', '<hostname>')
```

In addition, it is also possible to create custom location definitions, by using the following procedure to assign worker groups to a host:

```
call SYS.UPDATE_LANDSCAPE_CONFIGURATION( 'SET  
WORKERGROUPS', '<hostname>', '<name1> <name2> <name3>' )
```

Please note, if a host is assigned to several worker groups they must be specified separated by a space.

Note

- In `TABLE_PLACEMENT`, #1 can be used as affix to the location group name. This indicates that the partitions of all tables for which that rule apply will be located on a single host of that location group, that is, no distribution of several hosts takes place.
- If a location group (worker group) is not found in `M_LANDSCAPE_HOST_CONFIGURATION`, the system checks whether the location group has been defined by the HANA 1.0 configuration. This means whether an entry with the custom location name exists in section `[table_placement]` in `global.ini` file, assigning a comma separated list of volume IDs as value to it.

How Rules Are Applied

The `TABLE_PLACEMENT` table is read in such a way that a more specific rule supersedes a more generic one. A complete matrix of priorities is available in SAP Note 1908082 Table Placement Priorities.

For example, an entry with only schema applies to all tables of that schema; additional entries for that schema and specific group types overrule the more general rule.

Monitoring View

You can see the actual table placement settings per table by querying the system view `M_EFFECTIVE_TABLE_PLACEMENT`. You can see the valid location(s) according to the configuration and for each partitioning parameter the actual values and in the corresponding `_MATCH` columns the reason (matching rule) for those.

The information how a table is classified, can be reviewed in the monitoring view `SYS.TABLE_GROUPS`.

Related Information

[Dynamic Range Partitioning \[page 555\]](#)

[SAP Note 1908082](#)

6.8.10.3 Pre-defined Table Placement Scenarios

For specific applications, SAP provides recommendations regarding partitioning and table distribution configurations.

For the following scenarios SAP provides recommendations regarding table distribution configurations, for these scenarios SQL implementation scripts and detailed documentation is provided in SAP Notes.

SAP BW powered by SAP HANA

All required steps and recommended settings for SAP BW on HANA 2 are described in SAP Note 1908075 *BW on SAP HANA: Table placement and landscape redistribution*. This includes a zip file with documentation and SQL code to configure various scenarios covering a range of TABLE_PLACEMENT settings depending on the node size (TB per node) and the number of master and slave nodes.

SAP BW/4 HANA

All required steps and recommended settings for SAP BW/4 on HANA 2 are described in SAP Note 2334091 *BW/4HANA: Table Placement and Landscape Redistribution*. This includes a zip file with documentation and SQL code to configure various scenarios covering a range of TABLE_PLACEMENT settings depending on the node size (TB per node) and the number of master and slave nodes.

SAP S/4HANA

For details of scale-out options for SAP S/4HANA refer to SAP Note 2408419 *SAP S/4HANA - Multi-Node Support*. This note includes scripts and configuration settings as well as detailed documentation about table groups and migration.

SAP Business Suite Powered by SAP HANA and S/4HANA

SAP Note 1899817 *SAP Business Suite on SAP HANA database: table placement* includes configuration scripts to set up partitioning and distribution for Suite and S/4 HANA for various support package stack releases.

Further Resources

An introductory Frequently-Asked-Questions document is available on the SCN Community portal: *SAP BW/4HANA and SAP BW-on-HANA with SAP HANA Extension Nodes*

A scenario for SAP BPC (HANA 1 release) is described in SAP Note 2003863 *Enable BPC HANA table distribution*, this note includes table placement settings for BPC systems.

Related Information

[SAP BW/4HANA and SAP BW-on-HANA with SAP HANA Extension Nodes](#) 

[SAP Note 1908075](#) 

[SAP Note 2334091](#) 

[SAP Note 1899817](#) 

[SAP Note 2408419](#) 

6.8.10.4 Data Temperature: Extension Nodes

For scaled-out systems (SAP HANA native systems or in a SAP Business Warehouse) where a multi-temperature storage strategy is required, you can use the extension node feature to deploy a different type of host in the server landscape which is used exclusively for warm data.

Overview

The hardware sizing guidelines for data storage depend upon the type of data being stored. Hot data must be stored in SAP HANA memory for fast access, and the normal ratio of RAM to data storage in this case is 2:1, for example, a system with 8TB of RAM can store up to 4TB of hot data.

You can create an extension node in SAP HANA native systems or in SAP Business Warehouse. SAP BW offers powerful data modeling features which make it possible to separate warm data from hot data. Warm data must still be immediately available in the database but it can be stored on a separate node in the system landscape. This is referred to as an 'extension node'; it is used exclusively for warm data and because of this a more relaxed sizing ratio can be applied with less RAM in proportion to the volume of data. In this situation, a ratio of 1:1 may be possible so that a node with 2TB RAM could store 2TB of warm data. The economical advantages in this solution are not only the more efficient use of storage but also that expensive high-performance hardware may not be necessary for this node.

An extension node can be implemented by purchasing new hardware or by reconfiguring an existing node in the SAP HANA landscape. The configuration is based on a host sub role (worker group value) which identifies the special function of the node and a location value in the TABLE_PLACEMENT table. Scripts are available to configure a host with the worker group and location values required. Details are provided in the following two SAP Notes, firstly, for Business Warehouse and secondly a generic description for SAP HANA native scenarios:

- 2453736 - *How-To: Configuring SAP HANA for SAP BW Extension Node in SAP HANA 2.0 SP3*
- 2415279 - *How-To: Configuring SAP HANA for the SAP HANA Extension Node*

After modeling the data correctly to categorize the data according to temperature you can redistribute data using landscape management or data distribution optimizer so that objects that are classified for warm data will be moved to the extension node.

Configuration Using Worker Groups

The role of a host in an SAP HANA scale-out landscape is normally 'WORKER'. The extension node additionally has a sub role which is defined as a *Worker Groups* value. The name for this sub role must be defined as *worker_dt*; this is the name expected in Business Warehouse. The worker group is normally set up during installation and it can also be maintained in SAP HANA studio on the **► Landscape ► Hosts ►** tab page after clicking the *Configure Hosts for Failover Situation* button on the toolbar. For monitoring and administration the value is visible in the view SYS.M_LANDSCAPE_HOST_CONFIGURATION (WORKER_CONFIG_GROUPS).

The name of the worker group is also saved in the *Location* column in system view TABLE_PLACEMENT.

Note that the worker group name given to the extension node is also used by the following processes where it is important that the warm data is handled correctly:

- Backup and recovery to ensure that the right data is restored to this server
- SAP HANA system replication
- Standby nodes.

The use of sub roles and placing data in this way is not restricted to extension nodes, it can also be used more generically, for example, to pin schemas or applications to specific nodes of the landscape.

Tools Available for Redistribution

Data objects which already exist in SAP Business Warehouse can be reclassified and moved to the extension node using either the landscape redistribution tool in SAP HANA or the data distribution optimizer (DDO).

In SAP Business Warehouse data objects for warm data are mainly write-optimized data store objects (DSO), advanced DSOs (aDSO) or persistent staging area (PSA) tables. In the case of advanced DSOs which contain hot and warm data and are split by partitions, use the data distribution optimizer tool to redistribute the data.

Related Information

[Monitoring Host Status and Auto-Failover Configuration \[page 367\]](#)

[Redistributing Tables in a Scaleout SAP HANA System \[page 600\]](#)

[SAP Note 2343647](#)

[SAP Note 2453736](#)

[SAP Note 2415279](#)

6.8.11 Table Replication

In a scale-out system tables may be replicated to multiple hosts. This can help to reduce network traffic when, for example, slowly-changing master data often has to be joined with tables, or partitions of tables, that are located on other hosts.

Synchronous and Asynchronous Replication

Both asynchronous table replication (ATR) and synchronous table replication (STR) are supported. The following table summarizes the advantages and disadvantages of these two approaches.

With asynchronous table replication there can be differences between the source table and the replica table. This can cause issues because the application developer must decide which queries should see the replica

tables with the out-dated snapshot and then change the SQL or application. Synchronous table replication offers most of the benefits of ATR but is more transparent and no SQL or application changes are necessary.

Comparing the table replication types

Synchronous Table Replication	Asynchronous Table Replication
<p>Pro: Symmetric and thus easy to use.</p> <p>In STR, source and replica always have the same state. Therefore application developers do not need to be aware of existence of replicas. Queries can be routed to the source and replicas evenly and implicitly by the database.</p>	<p>Pro: Little overhead at the source node.</p> <p>Replicating updates with less overhead at the source transactions.</p>
<p>Con: There is a performance penalty but only to the write transactions commit operations (DML and read transactions are not affected).</p>	<p>Cons: not easy to use due to asymmetry between source and replica</p> <p>Replicas have different (possibly outdated) state than their source tables. This incurs difficulty in its usage model. That is, the source and its replica are not symmetric or equivalent to each other and the application developers should explicitly hint which queries are fine with such staleness.</p>

Note

In SAP HANA 2.0 a new version of synchronous table replication has been implemented (OSTR). Any replica tables which were created in earlier versions are no longer valid and must be recreated for the new HANA release. Refer to the *Synchronous Table Replication* topic for details.

Things to Consider

Before using table replication, consider the following aspects:

- Replicated tables consume main memory when they are loaded and disk space since each replica has its own independent persistence. Therefore an increased amount of main memory is needed. This needs to be considered for sizing.
- For replica creation, the configuration for table placement is taken into consideration.
- You cannot replicate column store tables that have history tables.

Table replication must only be used if these considerations are acceptable for your use case. Refer also to the topic which follows on general limitations.

Status Aware Routing

Status aware routing applies to both synchronous and asynchronous replication. This feature ensures that in a situation where replica tables exist, if replication is disabled then queries are automatically routed to the source table. For asynchronous replication it may be necessary to query a specific replica table using hints or routing, to do this replication must be enabled. See topic *Asynchronous Table Replication Operations* for details.

❁ Example

Creating Tables with Replicas on All Hosts

You can use the same syntax for both column store tables and row store tables.

SQL Command	Result
<pre>CREATE COLUMN TABLE MY_TABLE (I INT PRIMARY KEY) REPLICA AT ALL LOCATIONS</pre>	Creates a column store table with a replica on each available host.
<pre>ALTER TABLE MY_TABLE ADD REPLICA AT ALL LOCATIONS</pre>	Replicates an existing table on each available host
<pre>ALTER TABLE MY_TABLE2 DROP REPLICA AT ALL LOCATIONS</pre>	Drops all replicas

❁ Example

Creating Tables with Replicas on Specific Hosts

You can use the following commands to control the location of replicas.

SQL Command	Result
<pre>CREATE ROW TABLE MY_TABLE5 (I INT PRIMARY KEY) AT LOCATION '<master_node>;</pre>	Creates a new row store table with a replica on a specified host (for example, master)
<pre>ALTER TABLE MY_TABLE5 ADD REPLICA AT LOCATION '<first_slave_node>;</pre>	Replicates an existing row store table at the specified location
<pre>ALTER TABLE MY_TABLE5 MOVE REPLICA FROM '<first_slave_node>' TO '<second_slave_node>;</pre>	Moves an existing replica from one specified location to another
<pre>ALTER TABLE MY_TABLE3 DROP REPLICA AT LOCATION '<second_slave_node>;</pre>	Drops the replica from the specified host

Table Replication Consistency Checks

The following consistency checks are available for replicated tables in the column store.

To perform the general check, execute:

```
CALL CHECK_TABLE_CONSISTENCY('CHECK_REPLICATION', '<schema>', '<table>')
```

To perform a lightweight data check which ensures that all rows are replicated, execute:

```
CALL CHECK_TABLE_CONSISTENCY('CHECK_REPLICATION_DATA_LIGHTWEIGHT', '<schema>',  
'<table>')
```

To perform a full data check which ensures all replica hold the same data in all columns, execute:

```
CALL CHECK_TABLE_CONSISTENCY('CHECK_REPLICATION_DATA_FULL', '<schema>', '<table>')
```

The data checks can take a long time to run depending on the data volume.

Related Information

[Synchronous Table Replication \[page 597\]](#)

[Table Consistency Check \[page 509\]](#)

[Scaling SAP HANA \[page 1404\]](#)

[SAP Note 1819123](#)

[SAP Note 1908075](#)

6.8.11.1 Table Replication Limitations

General restrictions that apply to the use of table replication:

- Table replication works only within a single SAP HANA scale-out landscape. It cannot be used for replicating a table across two different SAP HANA landscapes or between SAP HANA and other DBMS instances.
- The following table types cannot be set as a replication table:
 - History table
 - Flexible table
 - Temporary table
 - Proxy table
 - Extended Storage table
 - Multistore table
 - System-versioned table
 - Tables with masked columns, tables with associations, tables with series data.
- Source tables can be distributed to multiple nodes but a source table and its replica cannot be located at the same node.
- Only one identical replica table (or table partition) can exist in a replica node.
- For partitioned tables, its source and replica should have identical partitioning specification.
- Write operations cannot be executed at the replica table. Such access will return an error.
- DDL operations cannot be executed directly at the replica table. Such access will return an error.
- DDL operations on a replication source table cannot be executed with a 'DDL auto-commit off' transaction. Such access will return an error.
- It is not allowed for a 'DDL auto-commit off' transaction to execute a write operation on a replication source table after a DDL operation on any table in the same transaction boundary.
- The binary import on the replicated table (source or replica table) is not allowed. You need to drop the replica tables and then import and re-create the replica tables again.
- Export from the replica table is not allowed. Export from the source table is allowed but the exported data contains only the source table's data. It does not contain any data of the replica.

6.8.11.2 Asynchronous Table Replication

Asynchronous table replication can help reduce workload on hosts by balancing load across replica tables on worker hosts in a distributed SAP HANA system.

Asynchronous replication has a number of key characteristics:

- **Table replication**
Only a selected list of tables can be set as replicated tables, this is different to system replication which replicates the entire database.
- **Asynchronous replication**
Write operations on those replicated tables are propagated to their replica tables asynchronously with almost no impact to the response time of source write transactions. That is, the write transaction is committed without waiting for its propagation to the replica.
- **Transactional replication**
Read queries routed to the replica may not see the up-to-date committed result by the nature of asynchronous replication. But, the cross-table transactional consistency is guaranteed by preserving the source transaction boundary and their commit order on log replay at the replica side.
- **Parallel log replay**
Although the read queries routed to the replica may see outdated data, the propagation delay is minimized by using parallel log replay at the replica side.

6.8.11.2.1 Configure Asynchronous Table Replication

To set up asynchronous table replication you create a replica schema, create replica tables, handle large column store tables and activate replication on the system.

Context

In the steps listed here SRC_SCHEMA, REP_SCHEMA, TAB1 and PART_TAB1 indicate source schema name, replica schema name, normal table name and partitioned table name respectively.

Procedure

1. Create a replica schema

```
CREATE SCHEMA REP_SCHEMA;
```

This creates the replica schema called REP_SCHEMA.

2. Create replica tables

You can choose the location of your replica tables using *Table Placement Rules* or you can *Set an Explicit Table Location*.

3. Handle large column store source tables.

If you need to create a replica of a large column store table (more than 2 GB) see *Optimize Replication Activation Time for Large Column Store Tables*.

4. Activate replication.

```
ALTER SYSTEM ENABLE ALL ASYNCHRONOUS TABLE REPLICAS;
```

Even after creating replica tables, replay is not activated without this activation command. We recommend first creating all of your necessary replicas and then activating them once.

Results

You have created and activated your table replicas.

You can check this with the following command:

```
SELECT * FROM M_ASYNCHRONOUS_TABLE_REPLICAS [WHERE SOURCE_SCHEMA_NAME = SRC_SCHEMA AND SOURCE_TABLE_NAME = 'TAB1'];
```

This will show all replica tables created for SRC_SCHEMA.TAB1.

Related Information

[Table Placement Rules for Replicas \[page 587\]](#)

[Set an Explicit Table Location \[page 589\]](#)

[Table Replication Limitations \[page 585\]](#)

6.8.11.2.1.1 Table Placement Rules for Replicas

The following SQL commands create table placement rules for replica schema, which can be used to create or add replica tables.

❖ Example

Create table placement rules

SQL Command	Result
<pre>ALTER SYSTEM ALTER configuration ('global.ini', 'SYSTEM') SET ('table_placement', 'repl_volumes')='6,7,8' WITH RECONFIGURE;</pre>	<p>This command creates a configuration parameter called 'repl_volumes' which will be used for the table placement rule.</p> <p>'6,7,8' means multiple nodes for replica host and each of them indicates a volume ID.</p>

SQL Command

```
ALTER SYSTEM ALTER TABLE PLACEMENT  
(SCHEMA_NAME => 'REP_SCHEMA') SET  
(LOCATION => 'repl_volumes');
```

Result

This command will create a table placement rule for 'REP_SCHEMA' which uses 'repl_volumes' as the configuration parameter. With this table placement rule, any of tables in REP_SCHEMA will be created at the locations mapped in 'repl_volumes'.

Note that replica table cannot be created on a master indexserver and it's not allowed to create multiple replicas of the same table on the same replica host.

❖ Example

Create or add a replica table using table placement rules

SQL Command

```
CREATE TABLE REP_SCHEMA.TAB1 LIKE  
SRC_SCHEMA.TAB1 ASYNCHRONOUS REPLICA;
```

Result

It creates the first replica table without assigning its location. The replica table will be created on one of replica hosts.

```
ALTER TABLE SRC_SCHEMA.TAB1 ADD  
ASYNCHRONOUS REPLICA;
```

This creates more than one replica for SRC_SCHEMA.TAB1

It creates the second replica table on another replica hosts without assigning its location. The second replica table will be created on one of the replica hosts which does not have the same replica table. If there are no more replica hosts to add a replica table to, an error will be returned. You can create the third, fourth, and so on, replica tables in the same manner.

```
CREATE TABLE REP_SCHEMA.PART_TAB1 LIKE  
SRC_SCHEMA.PART_TAB1 ASYNCHRONOUS  
REPLICA;
```

There's no difference between normal table and partitioned table. Please note that as a default, the replica's partitions will be created like source partitions' host distribution. Therefore the number of used hosts by the replica is equal to source table's one.

```
ALTER TABLE SRC_SCHEMA.PART_TAB1 ADD  
ASYNCHRONOUS REPLICA;
```

For partitioned table, it creates additional replica tables using the same procedure with the normal table. As a default, the additional replica partitions will be created on the replica hosts which don't already have the tables' replica partitions.

6.8.11.2.1.2 Set an Explicit Table Location

You can set an explicit table location with SQL commands.

❖ Example

SQL Command	Result
<pre>CREATE TABLE REP_SCHEMA.TAB1 LIKE SRC_SCHEMA.TAB1 ASYNCHRONOUS REPLICAS AT 'host1:port1';</pre>	<p>This creates the first replica table on the specified location 'host1:port1'.</p> <p>Note that replica table cannot be created on master index-server and it's not allowed to create multiple replica tables on the same replica node</p>
<pre>ALTER TABLE SRC_SCHEMA.TAB1 ADD ASYNCHRONOUS REPLICAS AT 'host2:port2'; ALTER TABLE SRC_SCHEMA.TAB1 ADD ASYNCHRONOUS REPLICAS AT 'host3:port3';</pre>	<p>Additional replica tables are created at the specified locations.</p>
<pre>CREATE TABLE REP_SCHEMA.PART_TAB1 LIKE SRC_SCHEMA.PART_TAB1 ASYNCHRONOUS REPLICAS AT ('host1:port1', 'host2:port2', ...);</pre>	<p>For partitioned table, this command creates the first replica table for a partitioned table. The replica partitions will be distributed on the specified nodes.</p>
<pre>ALTER TABLE SRC_SCHEMA.PART_TAB1 ADD ASYNCHRONOUS REPLICAS AT ('host3:port3', 'host4:port4', ...);</pre>	<p>For partitioned table, this command creates additional replica tables on other replica nodes. The replica partitions will be distributed on the specified nodes.</p>

6.8.11.2.2 Asynchronous Table Replication Operations

There are a number of operations you can perform on replica tables such as querying, adding, deactivating, dropping, and monitoring tables.

Querying Replica Tables

With asynchronous replication it may be necessary to query a specific replica table. To do this replication must be enabled, if it is disabled, all queries on replica tables are automatically re-routed to the source host and tables (this is called status aware routing).

If you submit a simple query to select data from a table which has multiple replica tables then one of the replica tables is automatically selected to service the query. However, it is also possible to use query hints to select a specific replica (identified by volume id) or to use the `result_lag` hint which only selects replica data if it is within an acceptable (specified) lag time. The following examples illustrate these methods:

Query Distribution

Using the following type of query one of the replica tables will be automatically selected:

```
SELECT * FROM REP_SCHEMA.TAB1;
```

Explicit Connection (Schema Mapping)

To access one specific replica table you can use the `route_to` hint to make an explicit connection to the location of the replica by including the volume ID number:

```
SELECT * FROM REP_SCHEMA.TAB1 with hint(route_to(4));
```

In this example '4' in the `route_to` hint identifies the volume id of the indexserver. If the specified volume has the replica table, it is selected to service the query.

You can use the following query to retrieve the volume id of a specific replica:

```
SELECT V.VOLUME_ID, C.SCHEMA_NAME, C.TABLE_NAME, C.PART_ID, C.RECORD_COUNT FROM  
M_VOLUMES V, M_CS_TABLES C  
WHERE V.HOST = C.HOST and V.PORT = C.PORT AND SCHEMA_NAME = 'REP_SCHEMA' AND  
TABLE_NAME LIKE '%TAB1%';
```

This example uses `M_CS_TABLES` to select a column table. Replace this with `M_RS_TABLES` to check for row tables.

Using Hints to Avoid Stale Data

You can query replica tables with the `result_lag` hint as shown in the following example.

```
SELECT * FROM SRC_SCHEMA.TAB1 WITH HINT(RESULT_LAG('hana_atr', [seconds]));
```

A preconfigured hint class exists for Asynchronous Table replication, which is called `hint_result_lag_hana_atr`. If the current lag time of the data on the replica is within the acceptable delay period (that is, if the current lag time is smaller than the stated value for the `[seconds]` parameter), then the query is executed on the replica. Otherwise the query is routed to the source table.

The seconds parameter is optional and if no value is entered on the SQL command line, the default value defined in configuration parameter `atr_default_lag_time` will be applied (configuration details are given in section *Performance: Using Hints to Query Data Snapshots*).

i Note

Note that if a query is submitted repeatedly and the staleness of the replica data in relation to the seconds parameter changes between query executions (from acceptable to unacceptable or from unacceptable to acceptable) a recompilation of the query would be triggered. Recompilation is triggered whenever the seconds parameter value of the hint is evaluated and the result causes a switch to a new data source.

It is important therefore in order to minimize the recompilation overhead to set an appropriate value for the seconds parameter in relation to how often the query is submitted and how frequently the data is refreshed.

Add More Replica Tables to an Active Asynchronous Table Replication System

You can create more replica tables in your existing ATR system and activate replication. You can activate table replication globally or for a specific named table as shown in the following examples:

```
ALTER SYSTEM ENABLE ALL ASYNCHRONOUS TABLE REPLICAS
```

This approach incurs a high-cost job if your system already has many replica tables or is actively replicating. In this case you are recommended to use the following command which requires global replication to be turned on:

```
ALTER TABLE SRC_SCHEMA.TAB2 ENABLE ASYNCHRONOUS REPLICA;
```

This example activates the replication operation for SRC_SCHEMA.TAB2. You can use the `disable` parameter instead of `enable` to deactivate replication. Note that during the table level replication activation phase, transactional consistency of the target replica table is not guaranteed.

Deactivate Replication

To deactivate the overall replication operation of all replication tables use:

```
ALTER SYSTEM DISABLE ALL ASYNCHRONOUS TABLE REPLICAS;
```

Its progress can be monitored by the following query:

```
SELECT * FROM M_TABLE_REPLICAS WHERE REPLICATION_STATUS != 'ENABLED'.
```

Drop Replica Tables

The following examples show how to drop replica tables. Note that if a source table is dropped its corresponding replica table is dropped as well.

This example drops REP_SCHEMA schema and all replica tables in the schema as well:

```
DROP SCHEMA REP_SCHEMA CASCADE;
```

These examples show dropping replica tables for SRC_SCHEMA.TAB1. Firstly at a specific named host and secondly at all locations:

```
ALTER TABLE SRC_SCHEMA.TAB1 DROP REPLICA AT '<replica host>:<replica port>';
```

```
ALTER TABLE SRC_SCHEMA.TAB1 DROP REPLICA AT ALL LOCATIONS;
```

Monitoring Replica Tables

Use the system view M_TABLE_REPLICAS to monitor replica tables. M_ASYNCHRONOUS_TABLE_REPLICAS is deprecated in SPS 12.

The field LAST_ERROR_CODE displays error codes. More detailed information will be described in field LAST_ERROR_MESSAGE.

You can look up the meaning of an error code in the system view M_ERROR_CODES. The error codes 2, 4 and 1025 are typically shown during replication and those are categorized as "ERR_GENERAL", "FATAL_OUT_OF_MEMORY" and "ERR_COM" respectively in M_ERROR_CODES.

Related Information

[Performance: Using Hints to Query Data Snapshots \[page 457\]](#)

6.8.11.2.3 Row to Column Table Replication

You can replicate data from row store tables to column store replicas, for mixed data types this may give optimal performance.

In a scale-out environment you can replicate data asynchronously from a row store source table to a column store replica table. Row store tables typically provide better performance for transactional (OLTP) workload in comparison to column store tables. Similarly, column store tables offer the best performance for analytics workload. Row to column table replication may therefore be an optimal replication configuration for mixed workload types to get the best performance from both types of table.

Row-to-Column Table Replication Operations

You can configure asynchronous Row-to-Column table replication with the following SQL commands:

- Create a row store source table:

```
CREATE ROW TABLE SRC_SCHEMA.TAB (A INT)
```

- Create a column store replica table. For the first replica creation you need to specify the 'COLUMN' keyword:

```
CREATE COLUMN TABLE REP_SCHEMA.TAB LIKE SRC.TBL AT '<host>:<port>'
```

or alter an existing table:

```
ALTER TABLE SRC_SCHEMA.TAB ADD ASYNCHRONOUS REPLICA AT '<host>:<port>'
```

All other operations are the same as general asynchronous table replication (such as table placement rule, activate replication, deactivate replication and drop replica tables).

Limitations

The following limitations apply:

- Only asynchronous mode is supported for Row-to-Column table replication. Synchronous mode is not supported because the replay performance of the column store does not fully catch up with the throughput of the row store source table.
- Binary object data is not supported, a row store source table which has an LOB type field cannot have a column store replica table.

6.8.11.2.4 Replicate Aging Tables

You can selectively replicate only the hot (current) partitions of aging tables, which means you can have the same benefit of the hot (current) partitions without increasing memory used for cold (old) partitions.

Procedure

1. Create an aging table.
 - a. You can create an aging table with the following SQL command:

Code Syntax

```
CREATE COLUMN TABLE SRC_SCHEMA.AGING_TABLE ( PK INT, TEMPERATURE
VARCHAR(8) default '00000000', PRIMARY KEY (PK) )
WITH PARAMETERS ('PARTITION_SPEC'='RANGE[TIME SELECTION] TEMPERATURE
00000000,*', 'LOCATION'=('host:port','host:port'))
```

- b. Promote a non-partitioned table into an aging table:

Code Syntax

```
ALTER TABLE SRC_SCHEMA.AGING_TABLE
WITH PARAMETERS('PARTITION_SPEC_ADD_RANGE_LEVEL'='RANGE[TIME SELECTION:
PAGED ATTRIBUTES, NO UNIQUE CHECK] TEMPERATURE 00000000')
```

- c. Promote a hash-partitioned table into an aging table

Code Syntax

```
ALTER TABLE SRC_SCHEMA.AGING TABLE
WITH PARAMETERS('PARTITION_SPEC_ADD_RANGE_LEVEL'='RANGE[TIME SELECTION:
PAGED ATTRIBUTES, NO UNIQUE CHECK] TEMPERATURE 00000000')
```

- Use RANGE for Range partitioning
- TIME SELECTION is the internal name for this Aging implementation
- PAGED ATTRIBUTES is an optional property that may be specified in order to use Paged Attributes for Cold partitions

- NO UNIQUE CHECK is an optional property that disables the unique check on Cold partitions
- TEMPERATURE is the VARCHAR(8) temperature column
- 00000000 is the identifier for the hot partition
- <ranges> shall be substituted with actual dates. For example, specify '20130101-20140101, 20140101-20150101'
- If an Aging table exists, use ADD PARTITION to create further Cold partitions.
For example, ALTER TABLE SRC_SCHEMA.AGING_TABLE ADD PARTITION 20000101 <= VALUES < 20020101

2. Optional: Enable Actual Only Replication.

In this release it is enabled by default.

3. Create a replica schema.

```
CREATE SCHEMA REP_SCHEMA
```

4. Activate replication.

Code Syntax

```
ALTER SYSTEM ENABLE ALL ASYNCHRONOUS TABLE REPLICAS;
```

This statement will activate all the other replicas except actual-only replication (actual-only replicas will be created in the next step). The actual-only replication should be enabled separately (in step 6) after all the other replicas are already enabled here.

5. Create replica tables.

a. Create replica table with Table Placement rule.

The following commands create table placement rules for a replica schema:

```
ALTER SYSTEM ALTER configuration ('global.ini', 'SYSTEM') SET ('table_placement', 'repl_volumes')='6,7,8' WITH RECONFIGURE
```

Here, *repl_volumes* is an alias name used to apply the table placement rule. '6,7,8' means multiple nodes are used as replica hosts and each number indicates the volume ID.

```
ALTER SYSTEM ALTER TABLE PLACEMENT (SCHEMA_NAME => 'REP_SCHEMA') SET (LOCATION => 'repl_volumes')
```

With this table placement rule, any of tables in REP_SCHEMA will be created at the locations mapped in "repl_volumes".

Note

Replica tables cannot be created on master indexserver and it is not allowed to create multiple replica tables on the same replica node.

To create the first replica table without assigning its location use the following SQL statement. The replica table will be created on one of replica nodes.

```
CREATE TABLE REP_SCHEMA.AGING_TABLE LIKE SRC_SCHEMA.AGING_TABLE ASYNCHRONOUS REPLICA
```

If you want to create more than one replica for SRC_SCHEMA, use:

```
ALTER TABLE SRC_SCHEMA.AGING_TABLE ADD ASYNCHRONOUS REPLICA
```

This creates the second replica table on other replica nodes without assigning its location. The second replica table will be created on one of replica nodes which does not have the same replica table. If there are no more replica nodes to add a replica table to, an error will be returned. You can create the third, the fourth, and more replica tables in the same manner.

- b. Create replica table with an explicit table location.

To create the first replica table on the specified location 'host:port'.

```
CREATE TABLE REP_SCHEMA.AGING_TABLE LIKE SRC_SCHEMA.AGING_TABLE ASYNCHRONOUS
REPLICA AT 'host:port'
```

Note

Replica tables cannot be created on master indexserver and it is not allowed to create multiple replica tables on the same replica node.

To create additional replica tables on the specified location:

```
ALTER TABLE SRC_SCHEMA.AGING_TABLE ADD ASYNCHRONOUS REPLICA AT 'host:port'
```

To create partitioned tables on more than one host use a comma separated list enclosed in parentheses as in the following examples:

```
CREATE TABLE REP_SCHEMA.AGING_TABLE LIKE SRC_SCHEMA.AGING_TABLE ASYNCHRONOUS
REPLICA AT ('host1:port1', 'host2:port2', ...)
```

```
ALTER TABLE SRC_SCHEMA.AGING_TABLE ADD ASYNCHRONOUS REPLICA AT
('host1:port1', 'host2:port2', ...)
```

6. Turn a partition on or off.

```
ALTER TABLE SRC_SCHEMA.AGING_TABLE [ENABLE/DISABLE] ASYNCHRONOUS REPLICA
PARTITION [logical partition id]
```

Only hot partition can be turned On/Off. If only Hot Partition is enabled, the others are not replicated.

7. Check Replica Tables

To view all replica tables created for SRC_SCHEMA.AGING_TABLE use:

```
SELECT * FROM M_ASYNCHRONOUS_TABLE_REPLICAS WHERE SOURCE_SCHEMA_NAME =
'SRC_SCHEMA' AND SOURCE_TABLE_NAME = 'AGING_TABLE'
```

8. Query on aging tables with a hint

You can read hot or cold data from an aging table using the following SQL suffix. The RANGE_RESTRICTION is a filter for the Range partitioning.

```
WITH RANGE_RESTRICTION('CURRENT') or WITH RANGE_RESTRICTION('DATE')
```

Use DATE in the format "yyyy-mm-dd". If you specify a date, it will always consider the hot partition as well. CURRENT is the hot partition.

Code Syntax

```
SELECT * FROM SRC_SCHEMA.AGING_TABLE WITH RANGE_RESTRICTION('CURRENT')
SELECT * FROM REP_SCHEMA.AGING_TABLE WITH RANGE_RESTRICTION('CURRENT')
SELECT * FROM SRC_SCHEMA.AGING_TABLE WITH RANGE_RESTRICTION('2000-01-01')
SELECT * FROM REP_SCHEMA.AGING_TABLE WITH RANGE_RESTRICTION('2000-01-01')
```

9. Deactivate Replication

```
ALTER SYSTEM DISABLE ALL ASYNCHRONOUS TABLE REPLICAS;
```

This command deactivates the overall replication operation of all replication tables.

You can monitor its progress using:

```
SELECT * FROM M_TABLE_REPLICAS WHERE REPLICATION_STATUS != 'ENABLED'.
```

You can turn off a specific table only using `ALTER TABLE SRC_SCHEMA.AGING_TABLE DISABLE ASYNCHRONOUS REPLICA.`

10. Drop Replica Tables

i Note

If a source table is dropped, its corresponding replica table is dropped as well.

```
DROP SCHEMA REP_SCHEMA CASCADE;
```

Drops REP_SCHEMA schema and all replica tables in the schema as well.

```
ALTER TABLE SRC_SCHEMA.AGING_TABLE DROP REPLICA AT ALL LOCATIONS
```

Drops all replica tables of the specified source table SRC_SCHEMA.AGING_TABLE.

```
ALTER TABLE SRC_SCHEMA.AGING_TABLE DROP REPLICA AT 'host:port'
```

Drops the replica located at '<replica host>:<replica port>'.

6.8.11.2.5 Query Aging Tables

An actual partition on a replica is only able to be accessed by using the CURRENT range restriction on a replica table. Otherwise, all queries are routed to a source table even though the queries are on a replica table.

Procedure

1. Access to actual partition(hot data) on replica

You can get hot data from an actual partition by using the CURRENT range restriction on a replica table.

```
SELECT * FROM REP_AGING_SCHEMA.AGING_TABLE WITH RANGE_RESTRICTION('CURRENT')
```

2. Access to both actual(hot data) and history partition(cold data) on replica

You can get hot and cold data from both actual and history partitions by using the DATE range restriction on a replica table. The query is routed to a source table. Even though the DATE range restriction indicates only hot data, the query is routed to a source table.

```
SELECT * FROM REP_AGING_SCHEMA.AGING_TABLE WITH RANGE_RESTRICTION('yyyy-mm-dd')
```

3. Access to replica without RANGE RESTRICTION

You can get data from all actual and historical partitions. The query is routed to a source table.

```
SELECT * FROM REP_AGING_SCHEMA.AGING_TABLE
```

6.8.11.3 Synchronous Table Replication

Synchronous table replication (STR) is a transparent solution which does not require any SQL or application changes. The table replication happens at commit time.

i Note

In SAP HANA 2.0 a new version of synchronous table replication has been implemented and replica tables which were created in earlier versions may no longer be valid. In this case an error message *Feature not supported* will be generated. If this error occurs the replica tables must be dropped and recreated using the ALTER TABLE commands for DROP and ADD:

```
ALTER TABLE {SCHEMA_NAME}.{TABLE_NAME} DROP REPLICA AT ALL LOCATIONS
```

```
ALTER TABLE {SCHEMA_NAME}.{TABLE_NAME} ADD [SYNCHRONOUS] REPLICA AT  
'replica_location'
```

Related Information

[Configure Synchronous Table Replication \[page 597\]](#)

[Asynchronous Table Replication Operations \[page 589\]](#)

6.8.11.3.1 Configure Synchronous Table Replication

You can configure synchronous replication using the SQL editor simply by adding replica tables.

Context

In the example commands below the following placeholders are used:

SRC_SCHEMA *	table schema name	REP_SCHEMA *	replica table schema name
SRC_TABLE *	table name	REP_TABLE	replica table name
SRC_PART_TABLE *	partitioned table name	REP_PART_TABLE	replica partitioned table name

*We assume that SRC_SCHEMA, REP_SCHEMA and SRC_TABLE and SRC_PART_TABLE already exist in your system.

Note that for synchronous replication no initial activation is necessary.

Procedure

Add Replica Tables

Normal Table	Partitioned Table
Explicit synchronous table creation	Explicit synchronous table creation
<pre>CREATE COLUMN TABLE REP_SCHEMA.REP_TABLE LIKE SRC_SCHEMA.SRC_TABLE SYNCHRONOUS REPLICA AT 'host:port'</pre>	<pre>CREATE COLUMN TABLE REP_SCHEMA.REP_PART_TABLE LIKE SRC_SCHEMA.SRC_PART_TABLE SYNCHRONOUS REPLICA AT 'host:port'</pre>
Implicit synchronous table creation	Implicit synchronous table creation
<pre>ALTER TABLE SRC_SCHEMA.SRC_TABLE ADD SYNCHRONOUS REPLICA AT 'host:port'</pre>	<pre>ALTER TABLE SRC_SCHEMA.SRC_PART_TABLE ADD SYNCHRONOUS REPLICA AT 'host:port';</pre>
This creates a replica table at the specified location 'host:port'. Note that it is not allowed to create multiple replica tables on the same replica node.	For a partitioned table, it creates a replica table for a partitioned table. The replica partitions cannot be distributed on the several nodes. They should be located on the same replica node.

Results

Activation is not necessary. You can check your replica tables by querying the M_TABLE_REPLICAS view.

This view shows general information on synchronous table replicas.

```
SELECT * FROM M_TABLE_REPLICAS WHERE SOURCE_TABLE_NAME = 'SRC_TABLE'
```

Related Information

[Table Replication Limitations \[page 585\]](#)

6.8.11.3.2 Operations for Synchronous Replication

There are a number of operations you can perform for synchronous replication such as activating or deactivating replication, and dropping replica tables.

Activate and Deactivate Replication

The following command deactivates all synchronous replication:

```
ALTER SYSTEM DISABLE ALL SYNCHRONOUS TABLE REPLICAS
```

The following command deactivates a specific synchronous table replica:

```
ALTER TABLE SRC_SCHEMA.SRC_TABLE DISABLE SYNCHRONOUS TABLE REPLICA
```

Activate Replication for all or specific synchronous tables again.

The following command activates all synchronous replication:

```
ALTER SYSTEM ENABLE ALL SYNCHRONOUS TABLE REPLICAS
```

The following command activates a specific synchronous table replica:

```
ALTER TABLE SRC_SCHEMA.SRC_TABLE ENABLE SYNCHRONOUS TABLE REPLICA
```

You can check a table's current replication status; this is shown in M_TABLE_REPLICAS.

Drop Replica Tables

Use the following commands to drop replica tables. Note that if a source table is dropped, its corresponding replica tables are all dropped as well.

The following command drops all replica tables of the specified source table "SRC_TABLE":

```
ALTER TABLE SRC_SCHEMA.SRC_TABLE DROP REPLICA AT ALL LOCATIONS
```

The following command drops one of replica tables which is located at "<replica host>:<replica port>":

```
SRC_TABLE'. ALTER TABLE SRC_SCHEMA.SRC_TABLE DROP REPLICA AT '<replica  
host>:<replica port>
```

Using synchronous replication you do not have to modify queries in any way to access replica tables. When you access a source table, either the source table or the replica table is automatically selected and the query is routed to the selected table in a round robin manner. So, just by adding replica tables, the read-only query workload can be load balanced without any SQL string or application problem change.

6.8.12 Redistributing Tables in a Scaleout SAP HANA System

In a scaleout SAP HANA system, tables and table partitions are assigned to an index server on a particular host when they are created. As the system evolves over time you may need to optimize the location of tables and partitions by running automatic table redistribution.

There are several occasions when tables or partitions of tables need to be moved to other servers, for example, the tables and partitions which grow fastest in size may need to be split and redistributed. Table redistribution aims to balance the workload across all hosts and optimize the location of tables and partitions so that tables which are often used together are located on the same node. Table redistribution is based on the table placement rules, these determine, for example, table sizes, partitioning threshold values, and preferred partition locations.

Although it is possible to move tables and table partitions manually from one host to another, this is not practical for large-scale redistribution of data. The table redistribution function offers a range of options to perform balancing, optimization and housekeeping tasks. Redistribution is a two stage process: the first stage is to generate a redistribution plan, this can be done iteratively and the distribution configuration can be modified and tweaked until the desired result is achieved; secondly the plan is executed.

You can run table redistribution from the SAP HANA administration tools, studio and cockpit, or from the SQL command line.

Related Information

[Managing Table Redistribution in SAP HANA Cockpit \[page 600\]](#)

[Managing Table Redistribution in SAP HANA Studio \[page 606\]](#)

[Table Redistribution Commands \[page 615\]](#)

[Table Placement \[page 574\]](#)

6.8.12.1 Managing Table Redistribution in SAP HANA Cockpit

Use the SAP HANA cockpit to manage table redistribution. You can view and save the current table distribution, automatically generate an optimized table distribution, re-run a previously executed plan, or restore a saved plan.

Context

In a scale-out system, tables and table partitions are distributed across multiple hosts. The location of the tables and partitions can affect performance when queries need to access several tables. You may want to redistribute the tables or partitions to better optimize for particular capabilities. Or you may want to add a new host to the scale-out system and therefore need to redistribute the tables so that some will reside on the new host.

Procedure

1. Open *Table Redistribution* in SAP HANA cockpit by clicking the *Manage Table Distribution* link on the system *Overview*.

Executed table distribution operations are displayed, including those that are:

- Running
 - Finished
 - Failed
 - Cancelled
2. Select a specific table distribution to drill-down to the table redistribution details.
 3. Choose the *Errors*, *All Statements*, or *Parameters* tabs to view specific information.
 4. Select a row to drill-down to details about steps.

Related Information

[Save Current Table Distribution \[page 601\]](#)

[Generate and Execute a Table Redistribution Plan \[page 602\]](#)

[View Table Distribution \[page 604\]](#)

[Rerun Table Distribution Plan \[page 605\]](#)

[Restore Saved Table Distribution Plan \[page 605\]](#)

6.8.12.1.1 Save Current Table Distribution

You can save a current table distribution as a distribution plan through the SAP HANA cockpit, provided that no table distribution operations are currently running.

Context

Changing how tables are distributed across the hosts of a distributed SAP HANA system is a critical operation. Therefore, before executing a redistribution operation, it is strongly recommended that you backup the landscape so that it can be restored if necessary.

Procedure

1. Open *Table Redistribution* in SAP HANA cockpit by clicking the *Manage Table Distribution* link on the system *Overview*.

2. In the table header, select the *Save Current Table Distribution* button.

→ Tip

This button is not available for selection if an executed table distribution is currently running.

3. Select *Save*.

The table distribution is saved as a distribution plan. You can rerun this plan at a later time if you wish.

Related Information

[Rerun Table Distribution Plan \[page 605\]](#)

[Restore Saved Table Distribution Plan \[page 605\]](#)

6.8.12.1.2 Generate and Execute a Table Redistribution Plan

Context

A table redistribution plan is temporary. Once the session is closed, the plan is removed. The most time consuming part is gathering all of the information about the existing landscape. You may wish to generate a plan, analyze the planned outcome, and regenerate a new plan using information that was previously gathered.

Procedure

1. Open *Table Redistribution* in SAP HANA cockpit by clicking the *Manage Table Distribution* link on the system *Overview*.
2. In the table header, select the *Generate Redistribution Plan* button.
3. Identify the goal of the table redistribution plan

Goal	Description
Balance table distribution	The load on a scale-out system changes over time with the usage of the system. This option generates a plan to move tables and partitions to their proper hosts if they are currently on invalid hosts according to the rules specified in the TABLE_PLACEMENT table. The plan will check whether a split or merge is necessary and calculates optimal positions for the parts and tables. All types of tables

Goal	Description
	and parts can be moved. However, only the tables that you have permission to view as catalog objects will be affected.
Check the number of partitions	<p>In a scale-out system, partitioned tables are distributed across different index servers. The location of the different partitions can be specified manually or determined by the database when the table is initially partitioned. Over time, this initial partitioning may no longer be optimal, for example, if a partition has grown significantly.</p> <p>This option evaluates whether or not partitioned tables need to be repartitioned. The plan will specify how partitioned tables will be repartitioned (split or merge) and how newly-created partitions will be distributed. Note that this is only relevant for column-store tables. System tables, temporary tables, and row-store tables are not considered.</p>
Redistribute tables after adding host(s)	After adding one or more worker hosts to a scaleout system, you may need to redistribute the tables across the active indexservers. This option checks whether new partitions can be created and generates a plan to move the tables and table partitions as necessary.
Check the correct location of tables and partitions	This option generates a plan to move tables and partitions to their proper hosts if they are on invalid hosts according to the rules specified in the TABLE_PLACEMENT table. Only the tables that you have permission to view as catalog objects will be affected.
Housekeeping	Some regular operations need to be done from time to time. This option allows you to perform various operations in the system, such as, optimize compressions, defrag, load table, merge delta. Only the tables that you have permission to view as catalog objects will be affected. Also, you must have appropriate privileges to perform specific housekeeping operations, such as delta merge.

4. Specify which tables to consider in the redistribution by setting any combination of:

- schema(s)
- table group(s)
- group type(s)
- group subtype(s)
- table name
- tables with or without LOB files
- loaded or unloaded tables
- filled or empty tables

- used or unused tables
5. Select the *Review* button.
The progress status displays on screen.
 6. After the plan has been generated successfully, select *Execute Plan*.

Related Information

[Save Current Table Distribution \[page 601\]](#)

6.8.12.1.3 View Table Distribution

To support the analysis and monitoring of performance issues in a distributed SAP HANA system, you can use the SAP HANA cockpit to see how tables are distributed across the hosts.

Context

In the case of partitioned tables, you can see how the individual partitions and sub-partitions are distributed, as well as detailed information about the physical distribution, for example, part ID, partition size, and so on.

Procedure

1. Open *Table Redistribution* in SAP HANA cockpit by clicking the *Manage Table Distribution* link on the system *Overview*.
2. In the table header, select the *View Table Distribution* button.
3. (Optional) Use the filtering options to refine the list of tables displayed according to table name and/or schema. You can display the list of tables grouped or not grouped, and indicate whether number of partition IDs and number of partitions are displayed.
4. Select *Go*.
A list of tables is displayed.

6.8.12.1.4 Rerun Table Distribution Plan

You can run a previously-executed table distribution plan through the SAP HANA cockpit.

Procedure

1. Open *Table Redistribution* in SAP HANA cockpit by clicking the *Manage Table Distribution* link on the system *Overview*.
2. Select a specific table distribution.
3. In the table header, select the *Rerun Plan* button.
4. In the dialog, confirm that you want to rerun the plan.

Related Information

[Save Current Table Distribution \[page 601\]](#)

6.8.12.1.5 Restore Saved Table Distribution Plan

You can restore a distribution plan through the SAP HANA cockpit.

Context

Changing how tables are distributed across the hosts of an SAP HANA system is a critical operation. You may need to restore the table distribution from a previous point in time.

Procedure

1. Open *Table Redistribution* in SAP HANA cockpit by clicking the *Manage Table Distribution* link on the system *Overview*.
2. Select a specific table distribution.
3. In the table header, select the *Restore Saved Distribution* button.
4. In the dialog, confirm that you want to restore the saved table distribution.

Related Information

[Save Current Table Distribution \[page 601\]](#)

6.8.12.2 Managing Table Redistribution in SAP HANA Studio

Administrators can use the table redistribution feature in the SAP HANA studio to create a plan for redistributing and repartitioning tables. The administrator can review the plan and execute it.

SAP HANA supports several redistribution operations that use complex algorithms as well as configurable table placement rules and redistribution parameters to evaluate the current distribution and determine a better distribution depending on the situation.

Redistribution operations are available to support the following situations:

- You are planning to remove a host from your system
- You have added a new host to your system
- You want to optimize current table distribution
- You want to optimize table partitioning

Data Distribution Optimizer

To plan, adjust and analyze landscape redistribution, you can also use the Data Distribution Optimizer. The Data Distribution Optimizer is an SAP HANA XS-based tool included in the SAP HANA Data Warehousing Foundation option. The Data Distribution Optimizer provides packaged tools for large scale SAP HANA use cases to support more efficient data management and distribution in an SAP HANA landscape. For more information, see the *SAP HANA Data Warehousing Foundation - Data Distribution Optimizer Administration Guide*.

Related Information

[Save Current Table Distribution \[page 607\]](#)

[Redistribute Tables Before Removing a Host \[page 608\]](#)

[Redistribute Tables After Adding a Host \[page 609\]](#)

[Restore Previous Table Distribution \[page 609\]](#)

[Optimize Table Distribution \[page 610\]](#)

[Optimize Table Partitioning \[page 611\]](#)

[Modify Table Distribution Manually \[page 612\]](#)

[Monitor Table Distribution \[page 614\]](#)

[Table Partitioning \[page 542\]](#)

[Table Replication \[page 582\]](#)

[Table Placement \[page 574\]](#)

6.8.12.2.1 Save Current Table Distribution

Changing how tables are distributed across the hosts of a distributed SAP HANA system is a critical operation. Therefore, before executing a redistribution operation, it is strongly recommended that you backup the landscape so that it can be restored if necessary.

Prerequisites

To be able to save the current table distribution, you must have the system privilege RESOURCE ADMIN and at least the object privilege ALTER and UPDATE for all schemas involved.

Procedure

1. In the Administration editor, choose **► Landscape ► Redistribution ►**.
2. Save the current table distribution by choosing *Save*.
The *Table Redistribution* dialog box appears.
3. Choose *Next*.
The system generates a redistribution plan that shows the distribution that will be saved.

i Note

Saving the current table operation involves the execution of a redistribution operation even though an actual redistribution of data does not take place.

4. Choose *Execute*.
The system saves the current table distribution. The associated redistribution operation appears in the list of executed operations.

Results

You can now restore the saved table distribution at later point in time if necessary.

6.8.12.2 Redistribute Tables Before Removing a Host

Before you can remove a host from your SAP HANA system, you must move the tables on the index server of the host in question to the index servers on the remaining hosts in the system.

Prerequisites

To be able to redistribute tables across the hosts in your system, you must have the system privilege RESOURCE ADMIN and at least the object privilege ALTER for all schemas involved. As redistributing data is a critical operation, it is also recommended that you have saved the current distribution so you can restore it if necessary.

Procedure

1. In the Administration editor, choose **► Landscape ► Hosts ▾**.
2. From the context menu of the host that you plan to remove, choose *Remove Host...*
3. In the *Remove Host* dialog box, choose *Yes*.
The system marks the host for removal and executes the required redistribution operation. This results in the data on the index server of the host being moved to the index servers of the remaining hosts in the system.

The redistribution operation appears in the list of executed operations on the *Redistribution* tab.

Results

You can remove the host.

Caution

After you remove the host from your system, you must perform a data backup to ensure that you can recover the database to a point in time after you removed the host.

Related Information

[Creating Backups \[page 1313\]](#)

6.8.12.2.3 Redistribute Tables After Adding a Host

After you have added a new worker host to your SAP HANA system, you need to redistribute the tables in the system to balance the memory footprint of the tables and to improve performance (load balancing).

Prerequisites

To be able redistribute tables across the hosts in your system, you must have the system privilege RESOURCE ADMIN and at least the object privilege ALTER and UPDATE for all schemas involved. As redistributing data is a critical operation, it is also recommended that you have saved the current distribution so you can restore it if necessary.

Procedure

1. In the Administration editor, choose **► Landscape ► Redistribution ►**.
2. In the *Redistribution Operations* area, select *Redistribute tables after adding host(s)* and choose *Execute*. The *Table Redistribution* dialog box appears.
3. Choose *Next*.
The system evaluates the current distribution of tables and generates a redistribution plan. This plan specifies which tables will be moved where.
4. Review the redistribution plan to ensure that you want to proceed and choose *Execute*.
The system redistributes the tables in your system across all available index servers. The associated redistribution operation appears in the list of executed operations.

Related Information

[Creating Backups \[page 1313\]](#)

6.8.12.2.4 Restore Previous Table Distribution

Changing how tables are distributed across the hosts of an SAP HANA system is a critical operation. You may need to restore the table distribution from a previous point in time.

Prerequisites

To be able to restore a previous table distribution, you must have the system privilege RESOURCE ADMIN and at least the object privilege ALTER for all schemas involved.

Procedure

1. In the Administration editor, choose ► *Landscape* ► *Redistribution* ►.
2. In the *Executed Operations* area, identify the operation that corresponds to the table distribution that you want to restore.
For example, you saved the table distribution at a particular point in time and you want to revert to this configuration.
3. Check the redistribution plan of the operation to ensure that you want to proceed.
To do this, select the operation and choose *Show Plan...*
4. Select the operation and choose *Restore*.
The system restores the selected table distribution. The associated redistribution operation appears in the list of executed operations.

6.8.12.2.5 Optimize Table Distribution

During production operation, you may discover that the initial assignment of tables and partitions to index servers is no longer optimal, for example, frequently joined tables are located on different servers. You can therefore trigger a redistribution operation that evaluates the current situation and determines how distribution can be improved.

Prerequisites

You must have the system privileges RESOURCE ADMIN and CATALOG READ, and at least the object privilege ALTER for all schemas involved.

Procedure

1. In the Administration editor, choose ► *Landscape* ► *Redistribution* ►.
2. In the *Redistribution Operations* area, select *Optimize table distribution* and choose *Execute*.
The *Table Redistribution* dialog box appears.
3. Optional: Enter the parameter `NO_SPLIT` if you do not want the operation to repartition tables.
4. Choose *Next*.
The system evaluates the current distribution of tables and partitions, and whether or not partitioned tables need to be repartitioned. A redistribution plan is subsequently generated, specifying which tables and partitions will be moved where, and how partitioned tables will be repartitioned and new partitions distributed.

i Note

The redistribution operation evaluates whether or not a partitioned table needs repartitioning based on its partitioning specification (that is, hash, round robin, range and so on). This is only relevant for column-store tables. System tables, temporary tables, and row-store tables are not considered.

5. Review the redistribution plan to ensure that you want to proceed and choose *Execute*. The system redistributes and repartitions the tables and partitions in your system.

i Note

Partitioning is a potentially expensive operation both in terms of time and memory consumption.

The associated redistribution operation appears in the list of executed operations.

6.8.12.2.6 Optimize Table Partitioning

In a distributed SAP HANA system, partitioned tables are distributed across different index servers. The location of the different partitions can be specified manually or determined by the database when the table is initially partitioned. Over time, this initial partitioning may no longer be optimal, for example, if a partition has grown significantly.

Prerequisites

To be able optimize partitioning, you must have the system privilege RESOURCE ADMIN and at least the object privilege ALTER for all schemas involved. As redistributing data is a critical operation, it is also recommended that you have saved the current distribution so you can restore it if necessary.

Procedure

1. In the Administration editor, choose ► *Landscape* ► *Redistribution* ►.
2. In the *Redistribution Operations* area, select *Optimize table partitioning* and choose *Execute*. The *Table Redistribution* dialog box appears.
3. Choose *Next*.

The system evaluates whether or not partitioned tables need to be repartitioned. A redistribution plan is subsequently generated, specifying how partitioned tables will be repartitioned and how newly-created partitions will be distributed.

i Note

The redistribution operation evaluates whether or not a partitioned table needs repartitioning based on its existing partitioning specification. This is only relevant for column-store tables. System tables, temporary tables, and row-store tables are not considered.

4. Review the redistribution plan to ensure that you want to proceed and choose *Execute*.
The system re-partitions the required tables and distributes the new partitions in your system.

i Note

Partitioning is a potentially expensive operation both in terms of time and memory consumption.

The associated redistribution operation appears in the list of executed operations.

Related Information

[Save Current Table Distribution \[page 607\]](#)

6.8.12.2.7 Modify Table Distribution Manually

In a distributed SAP HANA system, you can move individual tables or table partitions from the index server of one host to the index server of another.

Prerequisites

- You have the system privilege DATA ADMIN
- The target host has sufficient memory for the table(s) or partition(s).

Procedure

→ Recommendation

For a large-scale redistribution of data, it is recommended that you execute one of the available redistribution operations instead of modifying table distribution manually as described here. Redistribution operations use complex algorithms to evaluate the current distribution and determine a better distribution depending on the situation.

1. Open the table distribution editor by right-clicking the required entry in the *Systems* view and then choosing *Show Table Distribution*:
 - Catalog
 - Schema
 - Tables

A list of all tables is displayed.

i Note

For performance reasons, not all tables are displayed, but only the first 1,000. You can change this setting in the preferences of the SAP HANA studio under **► SAP HANA ► Runtime ► Catalog ►**. If more tables exist in the selected schema, a message is displayed.

2. Optional: Use the filtering options to refine the list of tables displayed according to table name and/or schema. If you want to see only those tables on a specific host or specific hosts, proceed as follows:
 - a. Open the table viewer by choosing *Configure Table...* from the context menu.
 - b. Move the hosts that you do not want to see from the *Visible Columns* column to the *Available Columns* column.
 - c. Close the table viewer.
 - d. In the table distribution editor, select the *Show only tables on selected hosts* checkbox.
3. To view the detailed distribution information of a partitioned table, select the table in the overview list. The information appears in the *Partition Details for <schema.table>* area.

i Note

You can only see table partition information for column-store tables as this is the only table type that can be partitioned.

4. To move a table to another host, proceed as follows:
 - a. Right-click the table in the overview list and choose *Move Table...*
 - b. Specify the host to which you want to move the table.If the target host has sufficient memory, the table is moved. The information in the table distribution editor is refreshed accordingly.
5. To move a table partition or sub-partition to another host, proceed as follows:
 - a. Right-click the partition or sub-partition in the *Table Partition Details* area and choose *Move Partitions...*
Note that you can select multiple partitions.
 - b. Specify the host to which you want to move the partition(s).If the target host has sufficient memory, the partitions are moved. The information in the table distribution editor is refreshed accordingly.

Related Information

[Table Partitioning \[page 542\]](#)

6.8.12.2.8 Monitor Table Distribution

To support the analysis and monitoring of performance issues in a distributed SAP HANA system, a table distribution editor is available in which you can see how tables are distributed across the hosts.

Context

In the case of partitioned tables, you can also see how the individual partitions and sub-partitions are distributed, as well as detailed information about the physical distribution, for example, part ID, partition size, and so on.

i Note

You can also see the detailed distribution information of an individual table by viewing its table definition (select *Table Definition* from the context menu when pointing at a table in the *Systems* view; partitioning information is shown on the *Runtime Information* tab).

Procedure

1. Open the *Table Distribution* editor by right-clicking any of the following entries in the *Systems* view and then choosing *Show Table Distribution*:
 - Catalog
 - Schema
 - Tables

A list of all tables is displayed.

i Note

For performance reasons, not all tables are displayed, but only the first 1,000. You can change this setting in the preferences of the SAP HANA studio under **▶ SAP HANA ▶ Runtime ▶ Catalog ▶**. If more tables exist in the selected schema, a message is displayed.

2. Optional: Use the filtering options to refine the list of tables displayed according to table name and/or schema. If you want to see only those tables on a specific host or specific hosts, proceed as follows:
 - a. Open the table viewer by choosing *Configure Table...* from the context menu.
 - b. Move the hosts that you do not want to see from the *Visible Columns* column to the *Available Columns* column.
 - c. Close the table viewer.
 - d. In the table distribution editor, select the *Show only tables on selected hosts* checkbox.
3. To view the detailed distribution information of a table, select the table in the overview list. The information appears in the *Partition Details of <schema.table>* area.

i Note

You can only see table partition information for column-store tables as this is the only table type that can be partitioned. A non-partitioned column-store table is considered a table with one partition.

Related Information

[Table Partitioning \[page 542\]](#)

[Redistributing Tables in a Scaleout SAP HANA System \[page 600\]](#)

6.8.12.3 Table Redistribution Commands

You can run table redistribution from the command line; this approach offers additional functionality including the option to modify at run time some of the configuration parameters which control redistribution.

Table redistribution is based on the table placement rules defined in the table TABLE_PLACEMENT, these determine, for example, table sizes, partitioning threshold values and preferred partition locations.

Redistribution is a two stage process: firstly to generate the plan and secondly to execute the plan; separate commands are used for each stage:

1. The plan generation command is a multi-purpose tool which requires an algorithm number as a parameter to determine which actions are executed. Depending on the algorithm selected, additional optional parameter values may also be available to give more control over the execution.
2. The plan execution command takes a single parameter which is the numeric plan id value. You can retrieve this value (REORG_ID) from the REORG_OVERVIEW system view - see System Views below.

The syntax for these commands is:

- `CALL REORG_GENERATE(<algorithm integer>, <optional parameter string>);`
- `CALL REORG_EXECUTE (<plan_id>)`

RESOURCE ADMIN and CATALOG READ privileges are required to call REORG_GENERATE(). The command only operates on tables and partitions which the executing user is allowed to see as catalog objects.

Generating the Plan: Algorithms and Options

The following table gives an overview of the most commonly-required algorithms and a summary of the options available for each one - see examples and details of the options which follow.

Table Redistribution Algorithms and Options

Num	Algorithm Name	Description
6	Balance landscape	<p>This function checks if tables in the landscape are placed on invalid servers according to the table placement rules, and checks if a split or merge is necessary in order to achieve optimal positions for the partitions and tables and to evenly distribute tables across the indexserver hosts.</p> <p>Options: SCHEMA_NAME TABLE_NAME GROUP_NAME GROUP_TYPE GROUP_SUBTYPE RECALC NO_PLAN NO_SPLIT SCOPE</p>
1	Add server	<p>Run this check after adding one or more index servers to the landscape. If new partitions can be created a plan will be generated to split the tables and move the new partitions to the newly added indexservers.</p> <p>Options: SCHEMA_NAME TABLE_NAME GROUP_NAME GROUP_TYPE GROUP_SUBTYPE RECALC NO_PLAN</p>
4	Save	Save current landscape setup. No optional parameter.
5	Restore	Restore a saved landscape setup. Enter the plan ID value as the optional parameter value.
7	Check number of partitions	This function checks if partitioned tables need to be repartitioned and creates a plan to split tables if the partitions exceed a configured row count threshold. No optional parameter.
14	Check table placement	<p>Check current landscape against table placement rules and (if necessary) provide a plan to move tables and partitions to the correct hosts.</p> <p>Additional Options: LEAVE_UNCHANGED_UNTOUCHED KEEP_VALID NO_SPLIT</p>
15	Rerun plan	<p>Rerun failed items from previously executed plans.</p> <p>Option: RERUN_ALL</p>
16	Housekeeping	<p>Perform housekeeping tasks. Additional privileges may be required for specific actions.</p> <p>Housekeeping Options: OPTIMIZE_COMPRESSION DEFRAG LOAD_TABLE MERGE_DELTA ALL</p>

Optional Parameters

The following table gives more details of the optional parameters which are available.

Option	Type	Detail
SCHEMA_NAME	String	Restrict redistribution to the named schema(s) - comma-separated list.
TABLE_NAME	String	Restrict redistribution to the named table(s) - comma-separated list.
GROUP_NAME	String	Restrict redistribution to the named group(s) - comma-separated list.

Option	Type	Detail
GROUP_TYPE	String	Restrict redistribution to the named group types(s) - comma-separated list.
GROUP_SUBTYPE	String	Restrict redistribution to the named sub types(s) - comma-separated list.
RECALC	True / False	If true then recalculate the landscape data of the last REORG_GENERATE run. This option works only if REORG_GENERATE has been called before within the same connection session. This parameter can be used to speed up plan generation with different parameters.
NO_PLAN	True / False	If true then the planning stage of generating the plan is skipped. This can be used with external tools when landscape data needs to be collected and a distribution must be calculated but might be modified.
SCOPE	Keyword	Use one or more of the following values (see example which follows) to restrict the scope of the redistribution to include only the named items specified by these keywords. The default value is 'ALL' so that all tables visible to the user are included in the redistribution. <ul style="list-style-type: none"> LOADED Tables which are loaded or partially loaded UNLOADED Tables which are not loaded FILLED Tables with a record count greater than 10 EMPTY Tables with a record count less than or equal to 10 USED Tables with a total execution count greater than 10 UNUSED Tables with a total execution count of less than or equal to 10 LOB Tables with LOB columns NOLOB Tables without LOB columns

Examples

Add server (algorithm 1)

With this algorithm you can use the optional filter parameters to, for example, restrict redistribution to specified schemas, tables, table groups and so on. The following example uses the SCHEMA_NAME option to generate a plan for all tables in schema SAPBWP.

```
CALL REORG_GENERATE (1, 'SCHEMA_NAME => SAPBWP')
```

Balance Landscape / Table (algorithm 6)

The following examples show the usage of optional parameters with this balancing algorithm:

If the options parameter string is left blank a plan is generated for all visible tables:

```
CALL REORG_GENERATE (6, '');
```

This example uses the GROUP_NAME option to generate a plan for all tables in three specified groups:

```
CALL REORG_GENERATE (6, 'GROUP_NAME=>TABLEGROUP1, TABLEGROUP2, TABLEGROUP3');
```

This example uses the `SCHEMA_NAME` option to generate a plan for all tables in schema `SAPBWP`:

```
CALL REORG_GENERATE (6, 'SCHEMA_NAME => SAPBWP');
```

This example show usage of the `SCOPE` option. The plan is restricted to only tables with a record count greater than 10 and which have no LOB columns.

```
CALL REORG_GENERATE (6, 'SCOPE=>FILLED,NOLOB');
```

System Views

The following system views show details of table redistribution. The last two views in this list show information about the most recent distribution operation; the details are deleted when the current connection to the database is closed.

- `REORG_OVERVIEW` Provides an overview of landscape redistributions.
- `REORG_STEPS` Shows details of the individual steps (items) of each plan.
- `REORG_PLAN` This view contains details of the last table redistribution plan generated with this database connection.
- `REORG_PLAN_INFOS` Showing details (as key-value pairs) of the last executed redistribution (algorithm value and parameters used).

Related Information

[Table Placement Rules \[page 577\]](#)

[Managing Table Redistribution in SAP HANA Cockpit \[page 600\]](#)

[Managing Table Redistribution in SAP HANA Studio \[page 606\]](#)

6.8.12.3.1 Configuration of Table Redistribution

The operation of table redistribution is managed by configuration parameters; some of these can be reconfigured for the current session at run-time.

Configuration

The precise operation of table redistribution is managed by sets of configuration parameters located in the service configuration file (typically `indexserver.ini`) section `table_redist`. Firstly, there are parameters controlling common behavior of table redistribution and secondly a set of parameters for fine tuning redistribution by applying weighting values.

These parameters can be defined by the administrator for the system generally, but for some values the system settings can be overridden by values submitted on the command line. This is done by entering the key-

value pair as a parameter value as shown in the following example. These settings can be combined with other parameters in a comma-separated string.

```
CALL REORG_GENERATE (6, 'BALANCE_BY_EXECUTION_COUNT=>True');
```

Settings passed in this way are only valid for the current user session.

The following tables list the configuration parameters which can be set:

Parameter	Data Type
ALL_MOVES_PHYSICAL	BOOLEAN
ASYMETRIC_CORRECTION_BY	INTEGER
DEBUG_BALANCE_DETAIL	BOOLEAN
DEBUG_EXPORT_DETAIL	BOOLEAN
DEBUG_SCORECARD_DETAIL	BOOLEAN
ENABLE_CONSISTENCY_CHECK	BOOLEAN
ENABLE_CONTENT_TABLE_REDIST	BOOLEAN
ENABLE_ENSURE_MOVES	BOOLEAN
ENABLE_MERGE	BOOLEAN
ENABLE_MULTI_STORE_TABLES	BOOLEAN
ENABLE_OPTIMIZE_COMPRESSION	BOOLEAN
ENABLE_RELOAD_TABLES	BOOLEAN
ENABLE_REPARTITIONING_WITH_GCD	BOOLEAN
ENABLE_ROW_STORE_TABLES	BOOLEAN
ENABLE_SYS_TABLE_REDIST	BOOLEAN
EXPORT_DATA	BOOLEAN
FORCE_PARTNUM_TO_SPLITRULE	BOOLEAN
MAX_PARTITIONS	INTEGER
MAX_PARTITIONS_LIMITED_BY_LOCATIONS	BOOLEAN
MAX_PLAN_ITERATIONS	INTEGER
MAX_ROWS_PER_PARTITION	INTEGER
MEM_LOAD_TRESHOLD	INTEGER

Parameter	Data Type
MIN_PLAN_ITERATIONS	INTEGER
MOVE_ROW_STORE	STRING
USE_GROUPS_FOR_DEPENDENCY	BOOLEAN
WORK_SEQUENCE	STRING
WORK_SEQUENCE_ASC	BOOLEAN
WORK_SEQUENCE_SORT	STRING

The following parameters (can also be set from the command line) control the optimization process. They are in pairs: firstly a boolean parameter and a corresponding numeric 'weight' parameter with the suffix '_WEIGHT', for example:

- BALANCE_BY_PARTNUM - boolean parameter: set to True to activate this pair
- BALANCE_BY_PARTNUM_WEIGHT - numeric value used as a factor (default = 1)

The parameters can be activated by setting the first one to **true** and then setting a numeric value for the weight parameter, this operates as a multiplication factor to calculate a priority value for this aspect of redistribution.

Parameter	Optimize by...
BALANCE_BY_PARTNUM	- by the number of partitions placed on the indexserver
BALANCE_BY_MEMUSE	- by memory usage based on size of table / partition
BALANCE_BY_TABLE_SIZE_HOSTED	
BALANCE_BY_ROWS	- by the number of rows of the tables/partitions
BALANCE_BY_WORKLOAD	
BALANCE_BY_EXECUTION_COUNT	- by an external provided execution count
BALANCE_BY_EXECUTION_TIME	- by an external provided execution time
BALANCE_BY_RANDOMIZER	
BALANCE_BY_TABLE_CLASSIFICATION	

Group Advisor

Table grouping identifies tables which are often used together so that during redistribution they can be located together on the same node in order to avoid cross-node communication in the landscape. The group advisor tool can be integrated into table redistribution to execute a preparation step which creates groups

automatically. The tool determines which tables are often used together by analyzing the current statement cache to find relationships between tables and it then (internally) makes recommendations about which tables should be located together. Table redistribution evaluates these recommendations before generating the plan.

When the plan is executed the table grouping information is persisted in the table `SYS.TABLE_GROUP` by setting a group name and setting the group type `'sap.join'`. Note that existing table groups which have been defined for other applications (Business Warehouse, for example) are not modified in any way by the Group Advisor.

To activate group advisor you must set the configuration parameter `USE_GROUP_ADVISOR` to `true` the grouping step will then be executed within plan generation.

6.9 Workload Management

The load on an SAP HANA system can be managed by selectively applying limitations and priorities to how resources (such as the CPU, the number of active threads and memory) are used. Settings can be applied globally or at the level of individual user sessions by using workload classes.

On an SAP HANA system there are many different types of workload due to the capabilities of the platform, from simple or complex statements to potentially long-running data loading jobs. These workload types must be balanced with the resources (CPU or memory) that are available to handle concurrent work. For simplicity we classify workload queries as transactional (OLTP) or analytic (OLAP). With a transactional query the typical response time is measured in milliseconds and these queries tend to be executed in a single thread. Analytic queries on the other hand tend to feature more complex operations using multiple threads during execution, this can lead to higher CPU usage and memory consumption compared with transactional queries.

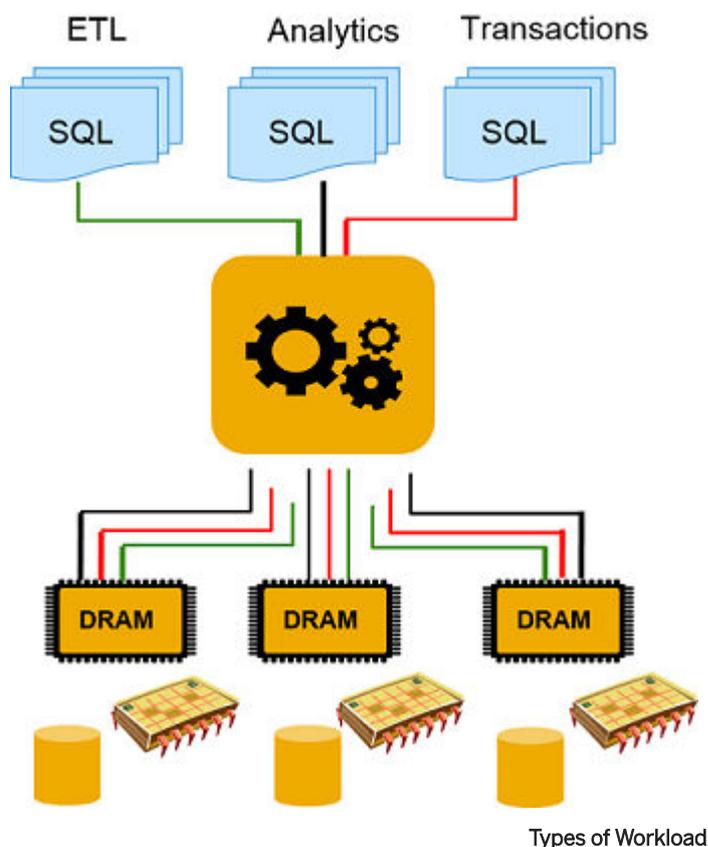
To manage the workload of your system aim to ensure that the database management system is running in an optimal way given the available resources. The goal is to maximize the overall system performance by balancing the demand for resources between the various workloads, not just to optimize for one particular type of operation. If you achieve this then requests will be carried out in a way that meets your performance expectations and you will be able to adapt to changing workloads over time. Besides optimizing for performance you can also optimize for robustness so that statement response times are more predictable.

6.9.1 Workload in the Context of SAP HANA

Workload in the context of SAP HANA can be described as a set of requests with common characteristics.

We can look at the details of a particular workload in a number of ways. We can look at the source of requests and determine if particular applications or application users generate a high workload for the system. We can examine what kinds of SQL statements are generated: are they simple or complex? Is there a prioritization of work done based on business importance, for example, does one part of the business need to have more access at peak times? We can also look at what kind of service level objectives the business has in terms of response times and throughput.

The following figure shows different types of workload such as Extract Transform and Load operations (used in data warehouses to load new data in batches from source system) as well as analytic and transactional operations:



When we discuss workload management we are really talking about stressing the system in terms of its resource utilization. The main resources we look at (shown in the above illustration) are CPU, memory, disk I/O, and network. In the context of SAP HANA, disk I/O comes into play for logging, for example, in an OLTP scenario many small transactions result in a high level of logging compared to analytic workloads (although SAP HANA tries to minimize this). With SAP HANA, network connections between nodes in a scale out system can be optimized as well, for example, statement routing is used to minimize network overhead as much as possible.

However, when we try to influence workload in a system, the main focus is on the available CPUs and memory being allocated and utilized. Mixed transactional and analytic workloads can, for example, compete for resources and at times require more resources than are readily available. If one request dominates there may be a queuing effect, meaning the next request may have to wait until the previous one is ready. Such situations need to be managed to minimize the impact on overall performance.

Related Information

[Persistent Data Storage in the SAP HANA Database \[page 462\]](#)

[Scaling SAP HANA \[page 1404\]](#)

6.9.1.1 Options for Managing Workload

Workload management can be configured at multiple levels: at the operating system-level, by using global initialization settings, and at the session level.

There are a number of things you can do to influence how workload is handled:

- Outside the SAP HANA system on the operating system level you can set the affinity of the available cores.
- You can apply static settings using parameters to configure execution, memory management and peak load situations.
- You can influence workload dynamically at system runtime by defining workload classes.

All of these options have default settings which are applied during the HANA installation. These general-purpose settings may provide you with perfectly acceptable performance in which case the workload management features described in this chapter may not be necessary. Before you begin with workload management, you should ensure that the system generally is well configured: that SQL statements are tuned, that in a distributed environment tables are optimally distributed, and that indexes have been defined as needed.

If you have specific workload management requirements the following table outlines a process of looking at ever more fine-grained controls that can be applied with regard to CPU, memory and execution priority.

Options for Controlling Workload Management

Area	Possible Actions
CPU Configure CPU at Operating System level	Settings related to <i>affinity</i> are available to bind server processes to specific CPU cores. Processes must be restarted before these changes become effective. For more information, see <i>Controlling CPU Consumption</i> .
CPU Thread Pools Configure CPU at HANA System level	Global <i>execution</i> settings are available to manage CPU thread pools and manage parallel execution (concurrency). For more information, see <i>Controlling Parallel Execution of SQL Statements</i> .
Memory Apply settings for memory management	Global <i>memorymanager</i> settings are available to apply limits to the resources allocated to expensive SQL statements. For more information, see <i>Setting a Memory Limit for SQL Statements</i> .
Admission Control Configuration options for peak load situations	Global <i>admission control</i> settings can be used to apply system capacity thresholds above which SQL statements can be either rejected or queued. For more information, see <i>Managing Peak Load (Admission Control)</i> .

Area	Possible Actions
Priority and Dynamic Workload Class Mapping Manage workload and workload priority using classes	<p>A more targeted approach to workload management is possible by setting up pre-configured classes which can be mapped to individual user sessions. You can, for example, map an application name or an application user to a specific workload class. Classes include the option to apply a workload priority value.</p> <p>You can set up classes using:</p> <ul style="list-style-type: none"> • SAP HANA cockpit • SQL commands <p>For more information, see <i>Managing Workload with Workload Classes</i>.</p>

At the end of this section is a set of scenarios giving details of different hardware configurations and different usage situations. For each scenario, suggestions are made about appropriate workload management options which could be used.

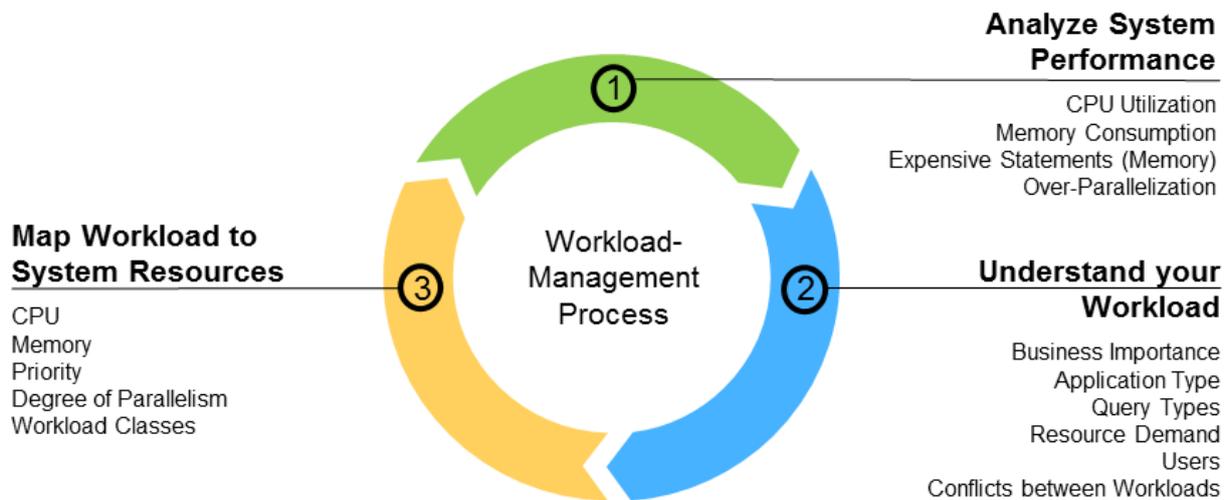
Related Information

- [Controlling CPU Consumption \[page 277\]](#)
- [Controlling Parallel Execution of SQL Statements \[page 630\]](#)
- [Setting a Memory Limit for SQL Statements \[page 633\]](#)
- [Managing Peak Load \(Admission Control\) \[page 636\]](#)
- [Managing Workload with Workload Classes \[page 640\]](#)
- [Example Workload Management Scenarios \[page 655\]](#)

6.9.1.2 Understand your Workload

Managing workload can be seen as an iterative three-part process: Analyze the current system performance, understand the nature of your workload and map your workload to the system resources.

There is no one single workload management configuration that fits all scenarios. Because workload management settings are highly workload dependent you must first understand your workload. The following figure shows an iterative process that you can use to understand and optimize how the system handles workload.



1. First you look at how the system is currently performing in terms of CPU usage and memory consumption. What kinds of workloads are running on the system, are there complex, long running queries that require lots of memory?
2. When you have a broad understanding of the activity in the system you can drill down in to the details such as business importance. Are statements being run that are strategic or analytic in nature compared to standard reporting that may not be so time critical. Can those statements be optimized to run more efficiently?
3. Then, when you have a deeper understanding of the system you have a number of ways to influence how it handles the workload. You can start to map the operations to available resources such as CPU and memory and determine the priority that requests get by, for example, using workload classes.

6.9.1.3 Analyzing System Performance

You can use system views to analyze how effectively your system is handling the current workload.

This section lists some of the most useful views available which you can use to analyze your workload and gives suggestions about actions that you might take to improve performance. Refer also to the scenarios section for more details of how these analysis results can help you to decide about which workload management options to apply.

Other performance analysis issues are described in the *SAP HANA Troubleshooting and Performance Analysis Guide*.

Analyzing SQL Statements

Use these views to analyze the performance of SQL statements:

- M_ACTIVE_STATEMENTS
- M_PREPARED_STATEMENTS

- M_EXPENSIVE_STATEMENTS

If these views indicate problems with statements you can use workload classes to tune the statements by limiting memory or parallelism.

Consider also the setting of any session variables (in M_SESSION_CONTEXT) which might have a negative impact on these statements. The following references provide more detailed information on this:

- SAP Note 2215929 *Using Client Info to set Session Variables and Workload Class settings* describes how client applications set session variables for dispatching workload classes.
- The *SAP HANA Developer Guide (Setting Session-Specific Client Information)*.

Analyzing CPU Activity

Use these views to analyze CPU activity:

- M_SERVICE_THREADS
- M_SERVICE_THREAD_SAMPLES
- M_EXPENSIVE_STATEMENTS.CPU_TIME (column)
- M_SERVICE_THREAD_CALLBACKS (stack frame information for service threads)
- M_JOBEXECUTORS (job executor statistics)

These views provide detailed information on the threads that are active in the context of a particular service and information about locks held by threads.

If these views show many threads for a single statement, and the general system load is high you can adjust the settings for the set of 'execution' ini-parameters as described in the topic *Controlling Parallel Execution*.

Related Information

[Controlling CPU Consumption \[page 277\]](#)

[Controlling Parallel Execution of SQL Statements \[page 630\]](#)

[Managing Workload with Workload Classes \[page 640\]](#)

[Example Workload Management Scenarios \[page 655\]](#)

[Managing and Monitoring the Performance of SAP HANA \[page 394\]](#)

[SAP Note 2215929](#)

6.9.2 Controlling CPU Consumption

If the physical hardware on a host is shared between several processes you can use CPU affinity settings to assign a set of logical cores to a specific SAP HANA process. These settings are coarse-grained and apply on the OS and process-level.

Prerequisites

Using this workload management option, we firstly analyze how the system CPUs are configured and then, based on the information returned, apply affinity settings in `daemon.ini` to bind specific processes to logical CPU cores. Processes must be restarted before the changes become effective. This approach applies primarily to the use cases of SAP HANA tenant databases and multiple SAP HANA instances on one server; you can use this, for example, to partition the CPU resources of the system by tenant database.

→ Tip

As an alternative to applying CPU affinity settings you can achieve similar performance gains by changing the parameter `[execution] max_concurrency` in the `global.ini` configuration file. This may be more convenient and does not require the system to be offline.

To make the changes described here you require access to the operating system of the SAP HANA instance to run the Linux `lscpu` command and you require the privilege INIFILE ADMIN.

Information about the SAP HANA system topology is also available from SAP HANA monitoring views as described in the following subsection *SAP HANA Monitoring Views for CPU Topology Details*.

Context

For Xen and VMware, the users in the VM guest system see what is configured in the VM host. So the quality of the reported information depends on the configuration of the VM guest. Therefore SAP cannot give any performance guarantees in this case.

Procedure

1. Firstly, to confirm the physical and logical details of your CPU architecture, analyze the system using the `lscpu` command. This command returns a listing of details of the system architecture. The table which follows gives a commentary on the most useful values based on an example system with 2 physical chips (sockets) each containing 8 physical cores. These are hyperthreaded to give a total of 32 logical cores.

#	Feature	Example Value
1	Architecture	x86_64
2	CPU op-mode(s)	32-bit, 64-bit
3	Byte Order	LittleEndian
4	CPUs	32
5	On-line CPU(s) list	0-31
6	Thread(s) per core	2
7	Core(s) per socket	8
8	Socket(s)	2
9	NUMA node(s)	2
21	NUMA node0 CPU(s)	0-7,16-23
22	NUMA node1 CPU(s)	8-15,24-31

- 4-5: This example server has 32 logical cores numbered 0 - 31
- 6-8: Logical cores ("threads") are assigned to physical cores. Where multiple threads are assigned to a single physical core this is referred to as 'hyperthreading'. In this example, there are 2 sockets, each socket contains 8 physical cores (total 16). Two logical cores are assigned to each physical core, thus, each core exposes two execution contexts for the independent and concurrent execution of two threads.
- 9: In this example there are 2 NUMA (Non-uniform memory access) nodes, one for each socket. Other systems may have multiple NUMA nodes per socket.
- 21-22: The 32 logical cores are numbered and specifically assigned to one of the two NUMA nodes.

i Note

Even on a system with 32 logical cores and two sockets the assignment of logical cores to physical CPUs and sockets can be different. It is important to collect the assignment in advance before making changes. A more detailed analysis is possible using the system commands described in the next step. These provide detailed information for each core including how CPU cores are grouped as siblings.

2. In addition to the `lscpu` command you can use the set of system commands in the `/sys/devices/system/cpu/` directory tree. For each logical core there is a numbered subdirectory beneath this node (`/cpu12/` in the following examples). The examples show how to retrieve this information and the table gives details of some of the most useful commands available:

❖ Example

```
cat /sys/devices/system/cpu/present
cat /sys/devices/system/cpu/cpu12/topology/thread_siblings_list
```

Command	Example Output	Commentary
present	0-15	The number of logical cores available for scheduling.
cpu12/topology/core_siblings_list	4-7, 12-15	The cores on the same socket.
cpu12/topology/thread_siblings_list	4, 12	The logical cores assigned to the same physical core (hyperthreading).
cpu12/topology/physical_package_id	1	The socket of the current core - in this case cpu12.

- Based on the results returned you can use the `affinity` setting to restrict CPU usage of SAP HANA processes to certain CPUs or ranges of CPUs. You can do this for the following servers: nameserver, indexserver, compileserver, preprocessor, and xsengine (each server has a section in the `daemon.ini` file). The examples and commentary below show the syntax for the ALTER SYSTEM CONFIGURATION commands required. The changed affinity settings only take effect after a restart of the affected SAP HANA processes.

❁ Example

To restrict the nameserver to two logical cores of the first CPU of socket 0 (see line 21 in the example above), use the following affinity setting:

```
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('nameserver', 'affinity') = '0,16'
```

❁ Example

To restrict the preprocessor and the compileserver to all remaining cores (that is, all except 0 and 16) on socket 0 (see line 21 in the example above), use the following affinity settings:

```
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('preprocessor', 'affinity') = '1-7,17-23'
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('compileserver', 'affinity') = '1-7,17-23'
```

❁ Example

To restrict the indexserver to all cores on socket 1 (see line 22 in the example above), use the following affinity settings:

```
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('indexserver', 'affinity') = '8-15,24-31'
```

❁ Example

To set the affinity for two tenant databases called DB1 and DB2 respectively in a tenant database setup, use the following affinity settings:

```
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('indexserver.DB1', 'affinity') = '1-7,17-23';
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
('indexserver.DB2', 'affinity') = '9-15,25-31';
```

Other Linux commands which are relevant here are `sched_setaffinity` and `numactl`:
`sched_setaffinity` limits the set of CPU cores available (by applying a CPU affinity mask) for execution of a specific process (this could be used, for example, to isolate tenants) and `numactl` controls NUMA policy for processes or shared memory.

Related Information

[Configuring Memory and CPU Usage for Tenant Databases \[page 266\]](#)

[SAP HANA Monitoring Views for CPU Topology Details \[page 630\]](#)

6.9.2.1 SAP HANA Monitoring Views for CPU Topology Details

A number of system views are available which you can use to retrieve details of the CPU configuration.

You can get a general overview of the system topology using the Linux `lscpu` command described above.

Information about the system topology is also available in the following system views:

M_HOST_INFORMATION provides host information such as machine and operating system configuration. Data in this view is stored in key-value pair format, the values are updated once per minute. For most keys the INIFILE ADMIN privilege is required to view the values. Select one or more key names for a specific host to retrieve the corresponding values:

```
select * from SYS.M_HOST_INFORMATION where key in  
( 'cpu_sockets', 'cpu_cores', 'cpu_threads' );
```

M_NUMA_RESOURCES provides information on overall resource availability for the system:

```
select HOST, NUMA_NODE_ID, NUMA_NODE_DISTANCES, MEMORY_SIZE from SYS.M_NUMA_NODES;
```

M_NUMA_NODES provides resource availability information on each NUMA node in the hardware topology, including inter-node distances and neighbor information.

```
select MAX_NUMA_NODE_COUNT, MAX_LOGICAL_CORE_COUNT from SYS.M_NUMA_RESOURCES;
```

Refer to the *SAP HANA SQL and System Views Reference Guide* for full details of these views.

6.9.3 Controlling Parallel Execution of SQL Statements

You can apply ini file settings to control the two thread pools `SqlExecutor` and `JobExecutor` that control the parallelism of statement execution.

⚠ Caution

The settings described here should only be modified when other tuning techniques like remodeling, repartitioning, and query tuning have been applied. Modifying the parallelism settings requires a thorough

understanding of the actual workload since they have impact on the overall system behavior. Modify the settings iteratively by testing each adjustment. For more information, see *Understand your Workload*.

On systems with highly concurrent workload, too much parallelism of single statements may lead to sub-optimal performance. Note also that partitioning tables influences the degree of parallelism for statement execution; in general, adding partitions tends to increase parallelism. You can use the parameters described in this section to adjust the CPU utilization in the system.

Two thread pools control the parallelism of the statement execution. Generally, target thread numbers applied to these pools are soft limits, meaning that additional available threads can be used if necessary and deleted when no longer required:

- **SqlExecutor**
This thread pool handles incoming client requests and executes simple statements. For each statement execution, an SqlExecutor thread from a thread pool processes the statement. For simple OLTP-like statements against column store as well as for most statements against row store, this will be the only type of thread involved. With OLTP we mean short running statements that consume relatively little resources, however, even OLTP-systems like SAP Business Suite may generate complex statements.
- **JobExecutor**
The JobExecutor is a job dispatching subsystem. Almost all remaining parallel tasks are dispatched to the JobExecutor and its associated JobWorker threads.
In addition to OLAP workload the JobExecutor also executes operations like table updates, backups, memory garbage collection, and savepoint writes.

You can set a limit for both SqlExecutor and JobExecutor to define the maximum number of threads. You can use this for example on a system where OLAP workload would normally consume too many CPU resources to apply a maximum value to the JobExecutor to reserve resources for OLTP workload.

⚠ Caution

Lowering the value of these parameters can have a drastic effect on the parallel processing of the servers and reduce the performance of the overall system. Adapt with caution by iteratively making modifications and testing. For more information, see *Understand your Workload* and *SAP Note 2222250 - FAQ SAP HANA Workload Management* which contains more details of the workload configuration parameters.

Parameters for SqlExecutor

The following SqlExecutor parameters are in the `sql` section of the `indexserver.ini` file.

`sql_executors` - sets a soft limit on the target number of logical cores for the SqlExecutor pool.

- This parameter sets the target number of threads that are immediately available to accept incoming requests. Additional threads will be created if needed and deleted if not needed any more.
- The parameter is initially not set (0); the default value is the number of logical cores in a system. As each thread allocates a particular amount of main memory for the stack, reducing the value of this parameter can help to avoid memory footprint.

`max_sql_executors` - sets a hard limit on the maximum number of logical cores that can be used.

- In normal operation new threads are created to handle incoming requests. If a limit is applied here, SAP HANA will reject new incoming requests with an error message if the limit is exceeded.

- The parameter is initially not set (0) so no limit is applied.

⚠ Caution

SAP HANA will not accept new incoming requests if the limit is exceeded. Use this parameter with extreme care.

Parameters for JobExecutor

The following JobExecutor parameters are in the `execution` section of the `global.ini` or `indexserver.ini`.

`max_concurrency` - sets the target number of logical cores for the JobExecutor pool.

- This parameter sets the size of the thread pool used by the JobExecutor used to parallelize execution of database operations. Additional threads will be created if needed and deleted if not needed any more. You can use this to limit resources available for JobExecutor threads, thereby saving capacity for SqlExecutors.
- The parameter is initially not set (0); the default value is the number of logical cores in a system. Especially on systems with at least 8 sockets consider setting this parameter to a reasonable value between the number of logical cores per CPU up to the overall number of logical cores in the system. In a system that supports tenant databases, a reasonable value is the number of cores divided by the number of tenant databases.

`max_concurrency_hint` - limits the number of logical cores for job workers even if more active job workers would be available.

- This parameter defines the number of jobs to create for an individual parallelized operation. The JobExecutor proposes the number of jobs to create for parallel processing based on the recent load on the system. Multiple parallelization steps may result in far more jobs being created for a statement (and hence higher concurrency) than this parameter.
- The default is 0 (no limit is applied but the hint value is never greater than the value for `max_concurrency`). On large systems (that is more than 4 sockets) setting this parameter to the number of logical cores of one socket may result in better performance but testing is necessary to confirm this.

`default_statement_concurrency_limit` - restricts the actual degree of parallel execution per connection within a statement.

- This parameter controls the maximum overall parallelism for a single database request. Set this to a reasonable value (a number of logical cores) between 1 and `max_concurrency` but greater or equal to the value set for `max_concurrency_hint`.
- The default setting is 0; no limit is applied.

User Priority Parameter

You can set a user-level priority value for all statements in the current connection; the range of possible values is from 0 to 9 (the default is 5). You can apply this from the SQL command line by setting the priority parameter as follows (USER ADMIN privileges are required):

```
ALTER USER MyUserId SET PARAMETERS PRIORITY '9';
```

Note that the user priority value is not effective in XSC developed applications. For XSC applications you can apply a priority value using workload classes.

The following examples show how to query the settings, firstly, to see the parameter settings currently applied for a selected user:

```
SELECT * FROM USER_PARAMETERS WHERE USER_NAME = 'SYSTEM';
```

Secondly, to see the priority currently applied for the current session:

```
SELECT PRIORITY FROM M_CONNECTIONS WHERE CONNECTION_ID = CURRENT_CONNECTION;
```

Related Information

[Understand your Workload \[page 624\]](#)

[Example Workload Management Scenarios \[page 655\]](#)

[SAP Note 2222250](#) 

6.9.4 Setting a Memory Limit for SQL Statements

You can set a statement memory limit to prevent single statements from consuming too much memory.

Prerequisites

To apply these settings you must have the system privilege INIFILE ADMIN.

For these options, `enable_tracking` and `memory_tracking` must first be enabled in the `global.ini` file. Additionally, `resource_tracking` must be enabled in this file if you wish to apply different settings for individual users (see Procedure below).

Context

You can protect an SAP HANA system from uncontrolled queries consuming excessive memory by limiting the amount of memory used by single statement executions per host. By default, there is no limit set on statement

memory usage, but if a limit is applied, statement executions that require more memory will be aborted when they reach the limit. To avoid canceling statements unnecessarily you can also apply a percentage threshold value which considers the current statement allocation as a proportion of the global memory currently available. Using this parameter, statements which have exceeded the hard-coded limit may still be executed if the memory allocated for the statement is within the percentage threshold. The percentage threshold setting is also effective for workload classes where a statement memory limit can also be defined.

You can also create exceptions to these limits for individual users (for example, to ensure an administrator is not prevented from doing a backup) by setting a different statement memory limit for each individual.

These limits only apply to single SQL statements, not the system as a whole. Tables which require much more memory than the limit applied here may be loaded into memory. The parameter `global_allocation_limit` limits the maximum memory allocation limit for the system as a whole.

You can view the (peak) memory consumption of a statement in `M_EXPENSIVE_STATEMENTS.MEMORY_SIZE`.

Procedure

1. Enable statement memory tracking.

In the `global.ini` file, expand the `resource_tracking` section and set the following parameters to **on**:

- `enable_tracking = on`
- `memory_tracking = on`

2. `statement_memory_limit` - defines the maximum memory allocation per statement in GB. The default value is 0 (no limit).

- In the `global.ini` file, expand the `memorymanager` section and locate the parameter. Set an integer value between 1 GB and the value of the global allocation limit. Values that are too small can block the system from performing critical tasks.
- When the statement memory limit is reached, a dump file is created with 'compositelimit_oom' in the name. The statement is aborted, but otherwise the system is not affected. By default only one dump file is written every 24 hours. If a second limit hits in that interval, no dump file is written. The interval can be configured in the `memorymanager` section of the `global.ini` file using the `oom_dump_time_delta` parameter, which sets the minimum time difference (in seconds) between two dumps of the same kind (and the same process).
- The value defined for this parameter can be overridden by the corresponding workload class property `STATEMENT_MEMORY_LIMIT`.

After setting this parameter, statements that exceed the limit you have set on a host are stopped by running out of memory.

3. `statement_memory_limit_threshold` - defines the maximum memory allocation per statement as a percentage of the global allocation limit. The default value is 0% (the `statement_memory_limit` is always respected).

- In the `global.ini` file, expand the `memorymanager` section and set the parameter as a percentage of the global allocation limit.
- This parameter provides a means of controlling when the `statement_memory_limit` is applied. If this parameter is set, when a statement is issued the system will determine if the amount of memory it consumes exceeds the defined percentage value of the overall `global_allocation_limit` parameter setting. The statement memory limit is only applied if the current SAP HANA memory

consumption exceeds this statement memory limit threshold as a percentage of the global allocation limit.

- This is a way of determining if a particular statement consumes an inordinate amount of memory compared to the overall system memory available. If so, to preserve memory for other tasks, the statement memory limit is applied and the statement fails with an exception.
 - Note that the value defined for this parameter also applies to the workload class property `STATEMENT_MEMORY_LIMIT`.
4. `total_statement_memory_limit` - limits the memory available to all statements running on the system.
- The value defined for this parameter cannot be overridden by the corresponding workload class property `TOTAL_STATEMENT_MEMORY_LIMIT`.
 - There is a corresponding parameter for use with system replication on an Active/Active (read enabled) secondary server. This is required to ensure that enough memory is always available for essential log shipping activity. See also `sr_total_statement_memory_limit` in section *Memory Management*.
5. To set a user-specific statement limit and exclude a user from the global limit use the `ALTER USER` statement as shown here to set the user parameter value:

```
ALTER USER myuser SET PARAMETER STATEMENT MEMORY LIMIT = <gb>
```

- Note that this user parameter-based approach to limiting memory for statements is not supported for cross-database queries, nor is it effective in XSC developed applications. In these cases you can apply memory limits using workload classes in the remote tenant database.
- If both a global and a user statement memory limit are set, the user-specific limit takes precedence, regardless of whether it is higher or lower than the global statement memory limit.
- If the user-specific statement memory limit is removed, the global limit takes effect for the user.
- The value of the parameter is shown in `USER_PARAMETERS` (like all other user parameters)

Note

Setting the statement memory limit to 0 will disable any statement memory limit for the user. Alternatively, to reset a user-specific limit use the `CLEAR` option:

```
ALTER USER myuser CLEAR PARAMETER STATEMENT MEMORY LIMIT
```

Results

The following example and scenarios show the effect of applying these settings:

Example showing statement memory parameters

Parameter	Value
Physical memory	128 GB

Parameter	Value
<code>global_allocation_limit</code>	The unit used by this parameter is MB. The default value is: 90% of the first 64 GB of available physical memory on the host plus 97% of each further GB; or, in the case of small physical memory, physical memory minus 1 GB.
<code>statement_memory_limit</code>	1 GB (the unit used by this parameter is GB.)
<code>statement_memory_limit_threshold</code>	60%

Scenario 1:

A statement allocates 2GB of memory and the current used memory size in SAP HANA is 50GB.

- $0,9 * 128\text{GB} = 115,2$ (global allocation limit)
- $0,6 * 115,2 = 69,12$ (threshold in GB)
- $50\text{ GB} < 69,12\text{ GB}$ (threshold not reached)

The statement is executed, even though it exceeds the 1GB `statement_memory_limit`.

Scenario 2:

A statement allocates 2GB and the current used memory size in SAP HANA is 70GB

- $70\text{ GB} > 69,12\text{ GB}$ (threshold is exceeded)

The statement is cancelled, as the threshold is exceeded, the `statement_memory_limit` is applied.

Related Information

[Change the Global Memory Allocation Limit \[page 478\]](#)

[Memory Management \[page 1165\]](#)

6.9.5 Managing Peak Load (Admission Control)

Use the admission control feature to apply processing limits and to decide how to handle new requests if the system is close to the point of saturation.

You can apply thresholds using configuration parameters to define an acceptable limit of activity in terms of the percentage of memory usage or percentage of CPU capacity.

Limits can be applied at two levels so that firstly new requests will be queued until adequate processing capacity is available or a timeout is reached, and secondly, a higher threshold can be defined to determine the maximum workload level above which new requests will be rejected. If requests have been queued, items in the queue are processed when the load on the system reduces below the threshold levels. If the queue exceeds a specified size or if items are queued for longer than a specified period of time they are rejected.

In the case of rejected requests an error message is returned to the client that the server is temporarily overloaded: `1038, 'ERR_SES_SERVER_BUSY', 'rejected as server is temporarily overloaded'`.

The load on the system is measured by background processes which gather a set of performance statistics covering available capacity for memory and CPU usage. The statistics are moderated by a configurable averaging factor (exponentially weighted moving average) to minimize volatility, and the moderated value is used in comparison with the threshold settings.

The admission control filtering process does not apply to all requests. In particular, requests that release resources will always be executed, for example, Commit, Rollback, Disconnect and so on. The filtering also depends on user privileges: administration requests from SESSION_ADMIN and WORKLOAD_ADMIN are always executed.

Admission control limitation with Active-Active (Read Only): admission control evaluates requests at the session layer as part of statement preparation before decoding the request packet from the client, whereas the decision about routing the request is made later in the SQL engine. This means that in an active-active setup if the admission control load threshold is exceeded on the primary then the incoming request is queued on the primary system. Statements which have been prepared and a decision to route the request to the secondary has already been made would directly connect to the secondary.

6.9.5.1 Configuring Admission Control

Threshold values for admission control to determine when requests are queued or rejected are defined as configuration parameters.

The admission control feature is enabled by default and the related threshold values and configurable parameters are available in the indexserver.ini file. A pair of settings is available for both memory and CPU which define firstly the queuing level (default value is 90%) and secondly the rejection level (not active by default). Two parameters are available to manage the statistics collection process by defining how frequently statistics are collected and setting the averaging factor which is used to moderate volatility. These parameters, in the *admission_control* section of the ini file, are summarized in the following table (see also Queue Management below):

Parameter	Default	Detail
enable	True	Enables or disables the admission control feature.
queue_cpu_threshold	90	The percentage of CPU usage above which requests will be queued. Queue details are available in the view M_ADMISSION_CONTROL_QUEUES.
queue_memory_threshold	90	The percentage of memory usage above which requests will be queued.
reject_cpu_threshold	0	The percentage of CPU usage above which requests will be rejected. The default value 0 means that no requests are rejected, but may be queued.
reject_memory_threshold	0	The percentage of memory usage above which requests will be rejected. The default value 0 means that no requests are rejected, but may be queued.
averaging_factor	70	This percentage value gives a weighting to the statistic averaging process: a low value has a strong moderating effect (but may not

Parameter	Default	Detail
		adequately reflect real CPU usage) and a value of 100% means that no averaging is performed, that is, only the current value for memory and CPU consumption is considered.
statistics_collection_interval	1000	Unit milliseconds. The statistics collection interval is set by default to 1000ms (1 second) which has a negligible effect on performance. Values from 100ms are supported. Statistics details are visible in the view M_ADMISSION_CONTROL_STATISTICS.

Events and Rejection Reasons

If statements are being rejected you may need to investigate why this is happening. Events related to admission control are logged and can be reviewed in the M_ADMISSION_CONTROL_EVENTS view. The key information items here are the event type (such as a statement was rejected or a statement was queued or dequeued) and the event reason which gives an explanatory text related to the type. Other details in this view include the length of time the statement was queued and the measured values for memory and CPU usage.

Two parameters are available to manage the event log in the *admission_control_events* section of the ini file:

Parameter	Default	Detail
queue_wait_time_threshold	100000	The length of time measured in microseconds for which a request must be queued above which it is included in the event log (default is one tenth of a second). If the parameter is set to 0 then events are not logged.
record_limit	1000000	The maximum record count permitted in the monitor of historical events.

Queue Management

If requests have been queued, items in the queue are processed when capacity becomes available. A background job continues to evaluate the load on the system in comparison to the thresholds and when the load is reduced enough queued requests are submitted in batches on an oldest-first basis.

The queue status of a request is visible in the M_CONNECTIONS view; the connection status value is set to *queuing* in column M_CONNECTIONS.CONNECTION_STATUS.

There are several configuration parameters (in the *admission_control* section of the ini file) to manage the queue and how the requests in the queue are released. You can apply a maximum queue size or a queue timeout value; if either of these limits are exceeded then requests which would otherwise be queued will be rejected. An interval parameter is available to determine how frequently to check the server load so that de-queueing can start, and a de-queue batch size setting is also available.

Parameter	Default	Detail
max_queue_size	10000	The maximum number of requests which can be queued. Requests above this number will be rejected.
dequeue_interval	1000	Unit: milliseconds. Use this parameter to set the frequency of the check to reevaluate the load in comparison to the thresholds. The default is 1000ms (1 second). This value is recommended to avoid overloading the system, though values from 100ms are supported.
dequeue_size	50	Use this parameter to set the de-queue batch size, that is, the number of queued items which are released together once the load is sufficiently reduced. This value can be between 1 and 9999 queued requests.
queue_timeout	600	Unit: seconds. Use this parameter to set the maximum length of time for which items can be queued. The default is 10 minutes. The minimum value which can be applied is 60 seconds, there is no maximum limit. Requests queued for this length of time will be rejected. Note that the timeout value applies to all entries in the queue. Any changes made to this configuration value will be applied to all entries in the existing queue.
queue_timeout_check_interval	10000	Unit: milliseconds. Use this parameter to determine how frequently to check if items have exceeded the queue timeout limit. The default is 10 seconds. The minimum value which can be applied is 100 milliseconds, there is no maximum limit.

i Note

If Admission Control has been configured and is active it takes precedence over any other timeout value which might have been applied. This means other timeouts which apply to a query (such as a query timeout) would not be effective until the query has been dequeued or rejected by the queue time out.

6.9.5.2 Use the Cockpit to Manage Admission Control

You can manage how new, incoming requests to run statements are rejected or queued based on memory and CPU statistics so as to manage peak load.

Prerequisites

CPU and Memory resource tracking are active. (You can use the [Configuration Manager](#) to set to `true` the two parameters `global.ini [resource tracking] enable_tracking` and `global.ini [resource tracking] memory_tracking`.)

Context

Admission control checks the current resource consumption within a system and in conjunction with the defined threshold values decides whether or not a new statement will be admitted during peak situations in the system.

Procedure

1. Open *Workload Classes* in SAP HANA cockpit by clicking the corresponding *DB Administration* link in the system *Overview*.
2. From the overflow menu above the table, select *Manage Admission Control*.
The cockpit displays *Workload Admission Control Settings*.
3. Ensure that there is a check-mark next to *Enable admission control so that requests can be rejected or queued based on CPU and memory resource tracking statistics*.
4. View or adjust the default settings for each of the parameters and thresholds.
5. Select *Save*.

Results

Any changes you make are saved as new settings in `global.ini [session_admission_control]`, or, in the case of admission control log management settings, in `global.ini [session_admission_control_events]`. You can view these new settings using the *Configuration Manager*.

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

6.9.6 Managing Workload with Workload Classes

You can manage workload in SAP HANA by creating workload classes and workload class mappings. Appropriate workload parameters are then dynamically applied to each client session.

You can classify workloads based on user and application context information and apply configured resource limitations (related to statement memory or thread limits) or a priority value. Workload classes allow SAP HANA to influence dynamic resource consumption on the session or statement level.

Workload class settings override other configuration settings (ini file values) which have been applied. Workload class settings also override user parameter settings which have been applied by the SQL command

ALTER USER, however, workload class settings only apply for the duration of the current session, whereas changes applied to the user persist. More detailed examples of precedence are given in a separate section.

To apply workload class settings client applications can submit client attribute values (session variables) in the interface connect string as one or more property-value pairs. The key values which can be used to work with workload classes are: database user, client, application name, application user and application type.

Based on this information the client is classified and mapped to a workload class. If it cannot be mapped it is assigned to the default workload class. The configuration parameters associated with the workload class are read and this sets the resource variable in the session or statement context.

The list of supported applications includes: HANA WebIDE (XS Classic), HANA Studio, ABAP applications, Lumira, and Crystal Reports. Full details of the session variables available in each supported client interface which can be passed in the connect string are given in SAP Note 2331857 *SAP HANA workload class support for SAP client applications*. Refer also to the *SAP HANA Developer Guide*, 'Setting Session-Specific Client Information'.

Managing workload classes requires the 'WORKLOAD ADMIN' privilege. Changes of workload classes or mappings will only be applied when a (connected) database client reconnects. In terms of the privileges of the executing user (DEFINER or INVOKER), the workload mapping is always determined on the basis of invoking user, regardless of if the user has definer or invoker privileges.

Users, classes and mappings are interrelated: if you drop a user in the SAP HANA database, all related workload classes are dropped and if you drop a workload class, the related mappings are also dropped.

i Note

In a scale-out environment workload classes are applied to the complete SAP HANA database but not to each single node.

Creating a Workload Class

You can use workload classes to set values for the properties listed here. Each property also has a default value which is applied if no class can be mapped or if no other value is defined. For all of the following parameters, although you can enter values including decimal fractions (such as 1.5GB) these numbers are rounded down and the whole number value is the effective value which is applied.

Parameter	Value
PRIORITY	To support better job scheduling, this property prioritizes statements in the current execution. Priority values of 0 (lowest priority) to 9 (highest) are available; the default value is 5.
STATEMENT THREAD LIMIT	To avoid excessive concurrent processing due to too many small jobs this property sets a limit on the number of parallel JobWorker threads per statement and process. The value can be set to a number between 1 and the number of logical cores. By default the value defined for the corresponding ini parameter <code>default_statement_concurrency_limit</code> is used. If that parameter is not set 0 (meaning no limit) is applied.

Parameter	Value
STATEMENT MEMORY LIMIT	To prevent a single statement execution from consuming too much memory this property sets a memory allocation limit in GB per statement. By default the value defined for the <code>statement_memory_limit</code> ini parameter is used. Note that if a percentage threshold value has been defined in the global <code>statement_memory_limit_threshold</code> parameter it is also effective for the workload class and may soften the effect of the statement memory limit.
TOTAL STATEMENT THREAD LIMIT	Similar to the STATEMENT THREAD LIMIT this property sets an aggregated thread limit which applies to all statements currently being executed within the workload class as a whole.
TOTAL STATEMENT MEMORY LIMIT	Similar to the STATEMENT MEMORY LIMIT this property sets an aggregated memory limit which applies to all statements currently being executed within the workload class as a whole.
STATEMENT TIMEOUT	This property applies a time limit (specified as a number of seconds) after which any query which has not completed execution will time out generating the following error: <code>ERR_API_TIMEOUT</code> .

Note

There are three ways of applying a query timeout value:

- At the command line (JDBC, SQLDBC) using a `querytimeout` instruction.
- Applying a time value (in seconds) in the `statement_timeout` configuration parameter (indexserver.ini file, session section).
- Applying a time value using workload classes.

If multiple values have been defined using different methods, precedence rules apply: the workload class takes precedence over the ini file, and, if a querytimeout value has been applied then the smallest (strictest) value which has been defined applies. See examples which follow.

Note

For thread and memory limits workload classes can contain either the statement-level properties or the aggregated total properties, not both. For the aggregated limits the full set of three properties must be defined: TOTAL STATEMENT THREAD LIMIT, TOTAL STATEMENT MEMORY LIMIT and PRIORITY.

Example

You can set values for one or more resource properties in a single SQL statement. This example creates a workload class called **MyWorkloadClass** with values for all three properties:

```
CREATE WORKLOAD CLASS "MyWorkloadClass" SET 'PRIORITY' = '3', 'STATEMENT MEMORY LIMIT' = '2' , 'STATEMENT THREAD LIMIT' = '20'
```

Examples of Precedence for Query Timeout

If multiple values have been defined using the different timeout methods available then precedence rules apply. Firstly, if a valid matching workload class value has been defined this takes precedence over the ini file setting. Secondly, if a querytimeout value has been applied then the smallest (strictest) valid value which has been defined applies. The following table shows some examples; in each case the values marked by an asterisk are the ones which apply.

QueryTimeout	25	25	25	25*
statement_timeout (ini)	10	10*	10*	10 (ignored)
STATEMENT TIMEOUT (Workload class)	20*	no match	no value	0 (disabled)

Creating a Workload Mapping

Mappings link workload classes to client sessions depending on the value of a specific client information property. The class with the most specific match is mapped to the database client.

The SAP HANA application sends client context information in the 'ClientInfo object'. This is a list of property-value pairs that an application can set in the client interface. You can change the running session-context of a connected database client using the SQL command 'ALTER SYSTEM ALTER SESSION SET', see also *Setting Session-Specific Client Information* in the *SAP HANA Developer Guide*.

The properties supported are listed here in prioritized order. The workload class with the greatest number of matching properties to the session variables passed from the client is applied. If two workload classes have the same number of matching properties then they are matched in the following prioritized order:

Field Name	Description
APPLICATION USER NAME	Name of the application user, usually the user logged into the application.
CLIENT	The client number is usually applied by SAP ABAP applications like SAP Business Suite / Business Warehouse.
APPLICATION COMPONENT NAME	Name of the application component. This value is used to identify sub-components of an application, such as CRM inside the SAP Business Suite.
APPLICATION COMPONENT TYPE	This value is used to provide coarse-grained properties of the workload generated by application components. In the future, SAP may document well-defined application component types to identify, for example, batch processing or interactive processing.
APPLICATION NAME	Name of the application
USER NAME	The name of the SAP HANA database user (the database the application is connected to).

Example

This example creates a workload mapping called **MyWorkloadMapping** which applies the values of the **MyWorkloadClass** class to all sessions where the application name value is **HDBStudio**:

```
CREATE WORKLOAD MAPPING "MyWorkloadMapping" WORKLOAD CLASS "MyWorkloadClass" SET
'APPLICATION NAME' = 'HDBStudio';
```

Refer also to Workload Management Statements in the *SAP HANA SQL and System Views Reference Guide* and refer to *Create a Workload Class Mapping* in this guide for details of maintaining workload classes in SAP HANA Cockpit.

Hints for Workload Classes

To give control over workload classes at run-time a *workload_class* hint is available. You can use this to apply more restrictive properties compared to the ones otherwise defined. For example, workload class 'YOUR_WORKLOAD_CLASS' applies the values: PRIORITY 5, THREAD 5, MEMORY 50GB. This is then overridden by the values defined in a new class, as a hint, to apply a higher priority value, a lower thread limit and a lower memory threshold:

```
SELECT * FROM T1 WITH HINT ( WORKLOAD_CLASS ("MY_WORKLOAD_CLASS") );
```

The complete hint is ignored if any of the new values are invalid or if they weaken the limits already applied. Refer to the *SAP HANA SQL and System Views Reference Guide* for full details.

Related Information

[Create a Workload Class Mapping \[page 651\]](#)

[SAP Note 2331857](#)

[SAP Note 2215929](#)

6.9.6.1 Managing Workload Classes

You can use system views to monitor, export, disable and view details of workload classes.

Monitoring

The following system views allow you to monitor workload classes and workload mappings:

- WORKLOAD_CLASSES
- WORKLOAD_MAPPINGS

In these system views the field `WORKLOAD_CLASS_NAME` shows the effective workload class used for the last execution of that statement:

- M_ACTIVE_STATEMENTS
- M_PREPARED_STATEMENTS
- M_EXPENSIVE_STATEMENTS (enable_tracking and memory_tracking must first be enabled in the global.ini file for this view)

- M_CONNECTIONS

If no workload class is applied then these views display the pseudo-workload class value "_SYS_DEFAULT".

You can also use queries such as the following examples to read data from these views:

Sample Code

```
-- get overview of available workload classes and workload class mappings
select wc.*, wm.workload_mapping_name, user_name, application_user_name,
application_name, client
from workload_classes wc, workload_mappings wm where wc.workload_class_name =
wm.workload_class_name;
```

Sample Code

```
-- get sum of used memory of all prepared statements grouped by workload
class which are executed in the last 10 minutes; requires memory tracking
select workload_class_name, sum(memory_size), count(*) statement_count,
count(distinct connection_id) as distinct_connection_count,
count(distinct application_name) as distinct_application_count,
count(distinct app_user) as distinct_applicationuser_count
from sys.m_expensive_statements
where add_seconds(start_time, 600) >= now() and memory_size >= 0
group by workload_class_name;
```

Sample Code

```
-- get information about priorities assigned to prepared statements executed
in the last 10 minutes
select workload_class_name, min(priority) min_priority, max(priority)
max_priority, count(*) statement_count,
count(distinct connection_id) as distinct_connection_count,
count(distinct application_name) as distinct_application_count,
count(distinct app_user) as distinct_applicationuser_count
from sys.m_expensive_statements
where add_seconds(start_time, 600) >= now()
group by workload_class_name;
```

Sample Code

```
-- collect workload related information for active statements
with job_count_per_statement as (select statement_id, count(0) num_active_jobs
from sys.m_service_threads
where statement_hash <> '' and is_active = 'TRUE'
group by statement_id, statement_hash)
select s.statement_id, s.statement_string, s.memory_size, s.duration_microsec,
s.application_source, s.application_name, s.app_user, s.db_user,
s.priority, s.statement_thread_limit, s.statement_memory_limit,
s.workload_class_name, st.num_active_jobs
from sys.m_expensive_statements s, job_count_per_statement st
where st.statement_id = s.statement_id;
```

Sample Code

```
-- collect workload related information for active statements
select s.statement_id, s.statement_string, s.memory_size, s.cpu_time,
s.application_source, s.application_name, s.app_user, s.db_user,
s.workload_class_name
```

```
from sys.m_expensive_statements s;
```

Sample Code

```
-- get information from system views
select * from sys.m_prepared_statements;
select * from sys.m_active_statements;
select * from sys.m_expensive_statements;
select * from m_service_threads;
select * from m_service_thread_samples;
```

Importing and Exporting Class Details

If you need to import and export workload classes, the normal SQL command for IMPORT will not work because workload classes do not belong to a schema (you cannot import into the SYS schema because it is maintained by the system). A script is available to support this functionality if you need to import a class. The script is typically shipped in the `exe/python_support` directory.

To use this you must first export the monitor views SYS.WORKLOAD_CLASSES and SYS.WORKLOAD_MAPPINGS to text (csv) format. You can then use the script to reimport the class:

1. Execute SQL EXPORT command:

```
EXPORT SYS.WORKLOAD_CLASSES, SYS.WORKLOAD_MAPPINGS AS CSV INTO '<PATH>' WITH REPLACE
```
2. Load CSV files using python script `importWorkloadClass.py` specifying the host as a parameter:

```
python importWorkloadClass.py --host='<host>' --SID='<SID>' --user='<DB-user>' --password='<PW>' <PATH>
```

Disabling and Enabling Workload Classes

After creating one or more workload classes it is also possible to disable them. This may be necessary, for example, for testing purposes.

1. Disable or enable a single named class:

```
ALTER WORKLOAD CLASS '<Class Name>' {enable | disable}
```
2. Disable all workload classes using the 'all' switch:

```
ALTER WORKLOAD CLASS all {enable | disable}
```

6.9.6.2 Workload Class Examples

Here we give examples to show how the workload management features interact together.

Workload class settings override other ini file configuration settings which have been applied and also override user parameter settings which have been applied by the SQL command ALTER USER, however, workload class settings only apply for the duration of the current session, whereas changes applied to the user persist. More

detailed examples of this and also how hints are applied to invoke a workload class at run-time are given in this section. The *SAP HANA SQL and System Views Reference* also includes examples of precedence including code samples.

i Note

For the examples shown in this section where memory limits are defined, the resource tracking parameters *enable_tracking* and *memory_tracking* in the global.ini file must be set to **true**.

Precedence of Workload Management Properties

In general, workload management settings are applied in the sequence of: workload class settings followed by user parameters followed by other ini file settings.

Precedence of Workload Class Values

The following table shows a scenario where a statement memory limit setting has been applied in the global ini file, restrictions have been applied at the level of the user profile and also in a workload class. In this basic example, if the workload class is mapped to a statement then the workload class settings take precedence and so for the duration of the current session the following values are effective: thread limit 10, memory limit 50 GB, priority 7.

Basic Scenario: Precedence of Workload Management Properties

	Global.ini	User Profile	Workload Class
Concurrency (job executor threads)	N/A	Thread limit 5	Thread limit 10
Statement Memory limit	50 GB	30 GB	50 GB
Priority	N/A	Priority 5	Priority 7

The following two variations on this scenario show the chain of precedence:

1. If the workload class thread limit is undefined the user profile value is effective: 5
2. If the workload class memory limit is undefined and the user profile memory limit is undefined then the global value is effective: 50 GB

Total Statement Values for Workload Classes

For workload classes using the aggregated properties (TOTAL STATEMENT THREAD LIMIT and TOTAL STATEMENT MEMORY LIMIT) the same principle of precedence applies as shown above. In this case all three values (thread limit, memory limit and priority) must all be defined and higher level settings would only be applied if the workload class was invalid for some reason.

Using Hints to specify a Workload Class

If a workload class is specified with an SQL statement as a hint at run-time then the settings of the matching hint class take precedence. In the following example, two workload classes are matched to the current execution request. The limitations specified by the hint class are successfully applied: thread limit 4, statement memory limit 30 GB, and priority 4.

Example showing use of hints

	Global.ini	User Profile	Workload Class	Hint Workload Class
Concurrency	N/A	Thread limit 5	Thread limit 5	Thread limit 4
Statement Memory limit	50 GB	30 GB	50 GB	30 GB
Priority	N/A	Priority 5	Priority 5	Priority 4

The following two variations on this scenario show the effects of additional conditions:

1. If, for example, the Concurrency (Thread limit) value of the hint workload class is left undefined then the two valid values of the hint class **are** applied and the Thread limit of the second matching workload class is applied (Thread limit 5 in the above example). If no valid value is found in a matching workload class then the value defined for the user profile is applied.
2. The hint can only be used to specify more restrictive values; in this scenario if the hint workload class memory limit was 70 GB then all values of the hint workload class would be ignored.

Related Information

[Managing Workload with Workload Classes \[page 640\]](#)

6.9.6.3 Managing Workload Classes in SAP HANA Cockpit

Several configuration options are available so that you can tailor workload classes in the SAP HANA database to your needs.

You can manage workload in SAP HANA by creating workload classes and workload class mappings. Workload classes and mappings are SQL object for workload management in SAP HANA. The goal of workload classes and mappings is to provide an easy way for administrators to regulate applications based on pre-defined mapping rules in order to avoid resource shortages with regard to CPU and memory consumption. Appropriate workload parameters are dynamically applied to each client session.

You can classify workloads based on user and application context information and apply configured resource limitations (for example, a statement memory limit). Workload classes allow SAP HANA to influence dynamic resource consumption on the session or statement level. When a request from an application arrives in SAP HANA, the corresponding workload class is determined based on the information given by the session context such as application name, application user name and database user name. Once the corresponding workload class is determined, the application request can have its resources limited according to the workload class definition.

Statement memory limits will not apply if memory tracking is inactive in SAP HANA cockpit. You can activate memory tracking in the Configuration settings.

Related Information

[Create a Workload Class \[page 650\]](#)

[Create a Workload Class Mapping \[page 651\]](#)

[Create User-Specific Parameters \[page 652\]](#)

[Apply Global Settings \[page 649\]](#)

[Disable or Enable a Workload Class \[page 653\]](#)

[Import Workload Classes \[page 654\]](#)

[Export Workload Classes \[page 655\]](#)

6.9.6.3.1 Apply Global Settings

You can apply global settings which are used as default values for workload classes. Enabling memory tracking allows you to also monitor the amount of memory used by single workload classes.

Context

Workload Classes lists existing workload classes and provides you with information about the workload handling of the database. You can create and edit workload classes and corresponding workload class mappings.

Procedure

1. Open *Workload Classes* in SAP HANA cockpit by clicking the corresponding *DB Administration* link in the system *Overview*.

The workload classes created in the database are listed. By default, workload classes are listed alphabetically. For each entry, you can see the execution priority, the statement memory limit, as well as the statement thread limit.

i Note

Not all of the columns listed below are visible by default. You can add and remove columns in the table personalization dialog, which you open by clicking the personalization icon in the table toolbar.

2. To monitor the memory consumption of workload classes, enable memory tracking using *Monitor Statements*.

Information about the memory consumption of workload classes is collected and displayed.

For more information about memory tracking and setting memory limits, see *Setting a Memory Limit for SQL Statements* in the *SAP HANA Administration Guide*.

3. To limit the memory consumption and number of threads per statement for the system globally select [Edit Global Limits](#) and specify the values for [Statement Memory Limit](#) and [Statement Thread Limit](#), then select [Save](#).
4. Click on a workload class entry in the list.

The mappings created for the workload class are listed, grouped, by default, by [Application User Name](#).

Related Information

[Managing Workload with Workload Classes \[page 640\]](#)

[Create a Workload Class \[page 650\]](#)

[Create a Workload Class Mapping \[page 651\]](#)

[Create User-Specific Parameters \[page 652\]](#)

[Disable or Enable a Workload Class \[page 653\]](#)

[Import Workload Classes \[page 654\]](#)

[Export Workload Classes \[page 655\]](#)

6.9.6.3.2 Create a Workload Class

You can create workload classes to manage the workload of the SAP HANA system.

Context

You can classify workloads based on user and application context information and apply configured resource limitations (for example, a statement memory limit). Workload classes allow SAP HANA to influence dynamic resource consumption . A workload class must contain at least one workload class mapping that specifies the workload based on user and application context information.

Procedure

1. Open [Workload Classes](#) in SAP HANA cockpit by clicking the corresponding [DB Administration](#) link in the system [Overview](#).
2. Select [Create](#). Specify the workload class details, then select [Create](#).

Field Name	Description
Workload Class Name	A name for the new workload class
Execution Priority	Priority, from 0 (lowest) to 9 (highest)

Field Name	Description
Statement Memory Limit	Maximum amount of memory the statement may use, in GB
Statement Thread Limit	Maximum number of parallel threads the statement may execute

3. You can also immediately create a mapping for the workload class by entering the mapping properties under *Mapping Details (Optional)*. Refer to *Creating a Workload Class Mapping* for details.

Results

The workload class is created and displayed in the list. If you have specified mapping properties, a mapping will also be created and assigned to the workload class.

Related Information

- [Create a Workload Class Mapping \[page 651\]](#)
- [Create a Workload Class Mapping \[page 651\]](#)
- [Create User-Specific Parameters \[page 652\]](#)
- [Disable or Enable a Workload Class \[page 653\]](#)
- [Import Workload Classes \[page 654\]](#)
- [Export Workload Classes \[page 655\]](#)

6.9.6.3.3 Create a Workload Class Mapping

Mappings link workload classes to client sessions depending on the value of a specific client information property. A workload class must contain at least one workload class mapping that specifies the workload based on user and application context information.

Procedure

1. Open *Workload Classes* in SAP HANA cockpit by clicking the corresponding *DB Administration* link in the system *Overview*.
2. Find the workload class to which you want to add a workload class mapping. Open the workload class by clicking on its entry in the list.
3. Workload class mappings are identified here on screen by a *Name* value. Select *Create*. Specify the mapping properties, then select *Create*.

The workload class with the greatest number of matching properties to the session variables passed from the client is applied. If two workload mappings have the same number of matching properties then they are

matched in the prioritized order as listed in the table: ► *application user name* ► *client* ► *application component name* ► *application component type* ► *application name* ► *user name* ». For example, a mapping where the application user matches takes precedence over a mapping where the database user matches (assuming an equal number of matching properties). SQL examples for this functionality are given in the section *Managing Workload with Workload Classes*.

Field Name	Description
APPLICATION USER NAME	Name of the user logged in to the application.
CLIENT	ABAP client number (Mandant). For example, 000.
APPLICATION COMPONENT NAME	Name of the application component. For example, /SSB/ ALERT_NOTIFICATION_REPORT.
APPLICATION COMPONENT TYPE	Name of the component type. For example, UPD.
APPLICATION NAME	Name of the application. For example, HDBStudio.
USER NAME	Name of the database user. For example, SYSTEM.

Results

The workload class mapping is created and displayed in the list.

Related Information

[Managing Workload with Workload Classes \[page 640\]](#)

[Create User-Specific Parameters \[page 652\]](#)

[Disable or Enable a Workload Class \[page 653\]](#)

[Import Workload Classes \[page 654\]](#)

[Export Workload Classes \[page 655\]](#)

6.9.6.3.4 Create User-Specific Parameters

User-specific parameters can be created for workload classes.

Context

You can set the execution priority and the statement memory limit for each database user individually. These settings will apply to all workload class mappings created for a given database user.

Procedure

1. Open [Workload Classes](#) in SAP HANA cockpit by clicking the corresponding [DB Administration](#) link in the system [Overview](#).
2. Select [User-Specific Parameters](#).
3. Select [Create](#). Specify the user-specific parameters, then select [Save](#).

Field Name	Description
Database User Name	Name of the database user
Execution Priority	Execution priority
Statement Memory Limit	Maximum amount of memory used to execute the statement

Results

The user-specific parameters are created and displayed in the list.

6.9.6.3.5 Disable or Enable a Workload Class

You can disable or enable workload classes.

Context

After creating one or more workload classes, you can disable them. This may be necessary for testing purposes. You can also enable workload classes that have been previously disabled.

Procedure

1. Open [Workload Classes](#) in SAP HANA cockpit by clicking the corresponding [DB Administration](#) link in the system [Overview](#).
2. To disable workload classes:
 - a. From the overflow menu above the table, select [Disable](#).
 - b. Select one or more workload classes.
 - c. Select [OK](#).

The workload classes are disabled.

3. To enable workload classes:

- a. Select *Enable*.
- b. Select one or more workload classes.
- c. Select *OK*.

The workload classes are enabled.

6.9.6.3.6 Import Workload Classes

Workload classes can be imported from another system, as in the case of going from a test system to a production system.

Prerequisites

You have a file containing workload classes exported from another system.

Procedure

1. Open *Workload Classes* in SAP HANA cockpit by clicking the corresponding *DB Administration* link in the system *Overview*.
2. From the overflow menu above the table, select *Import*.
3. Follow the prompts in the dialog to choose from where to retrieve the file containing previously exported workload class definitions.

Related Information

[Export Workload Classes \[page 655\]](#)

6.9.6.3.7 Export Workload Classes

Workload classes can be exported in preparation for importing them into another system, as in the case of going from a test system to a production system.

Procedure

1. Open *Workload Classes* in SAP HANA cockpit by clicking the corresponding *DB Administration* link in the system *Overview*.
2. Select one or more workload classes.
3. From the overflow menu above the table, select *Export*.
4. Follow the prompts in the dialog to specify where to save the file.

Related Information

[Import Workload Classes \[page 654\]](#)

6.9.7 Example Workload Management Scenarios

Here, we give a number of scenarios to illustrate how workload management settings can be applied for systems of different sizes, different workload types (analytics and transactional) and different usage scenarios (a system which is optimized for robustness as opposed to high performance).

i Note

All settings are tuning parameters and must be tested and validated before being used in production. See the process description in *Understand Your Workload Management*.

System Details

The scenarios which follow are based on the system specifications given here: firstly describing hardware resources and secondly the workload types which the system is expected to handle.

System Types 1: Small and Large Hardware Configurations

	Small (maximum 4 sockets)	Large
Sockets (processors)	2	16

	Small (maximum 4 sockets)	Large
Physical cores per socket	8	2 x 15 = 30
Logical cores (threads)	32	16 x 30 = 480
Memory	64 GB	256GB

Note that related to memory resources, the setting for global memory allocation limit is very important. By default, it is calculated as 90% of the first 64 GB of available physical memory on the host plus 97% of each further GB; or, in the case of small physical memory, physical memory minus 1 GB.

Secondly, we give details of two contrasting types of workload: pure analytics processing and mixed transactions, to show how the system can be configured to handle these different situations:

System Types 2: Workload and Processing Types

	Analytics Workload	Mixed Workload
Installed Applications	SAP Business Warehouse (or similar analytic application)	SAP Business Suite (OLTP) and Smart Business Apps (OLAP)
Workload type	Only OLAP.	Mixed OLAP and OLTP.
Processing characteristics	CPU and Memory intensive - long transaction times (+ 1 second)	Both applications have short-running statements (milliseconds or microseconds) where response time is critical as well as long-running CPU and memory-intensive OLAP statements.
Data	Bulk loading	
Concurrent queries, Concurrent users.	Few (10-100)	Many (> 1000)

Scenario 1: Mixed Workload (OLAP and OLTP) Optimized for Robustness

In the first scenario, the focus is on achieving robust statement execution time and high availability of the system.

Small System with Mixed Workload Optimized for Robustness

Tuning Option	Example Setting
<code>statement_memory_limit</code>	Start testing with this parameter set to 25% of the global allocation limit.
<code>default_statement_concurrency_limit</code>	Assign 33% of the logical cores on the system.

Tuning Option	Example Setting
<code>max_concurrency_hint, num_cores</code>	For systems with at least 80 logical cores consider <code>max_concurrency_hint</code> and <code>num_cores</code> equal to the number of logical cores per socket.
Workload Classes	Fine-grained control is possible with workload classes. In a mixed workload scenario, the OLTP load could be configured with a higher priority than OLAP.

Scenario 2: Mixed Workload (OLAP and OLTP) Optimized for Performance

For the second scenario we take exactly the same system as scenario 1 but we optimize for performance instead of robustness by relaxing some of the settings:

Small System with Mixed Workload Optimized for Performance

Tuning Option	Example Setting
<code>max_concurrency_hint, num_cores</code>	For systems with at least 80 logical cores consider <code>max_concurrency_hint</code> and <code>num_cores</code> equal to the number of logical cores per socket.
Workload Classes	Consider assigning higher priority to OLTP statements than to OLAP statements by using workload classes based on the application.

Scenario 3: Analytics

Pure OLAP scenarios tend to be more performance oriented, hence, we remove the limitations on concurrency here (leading to a best-effort approach). On the other hand, to avoid out-of-memory situations, we keep the memory limits.

Small Analytic System

Tuning Option	Example Setting
<code>statement_memory_limit</code>	Start testing with this parameter set to 25% of the global allocation limit.
<code>statement_memory_limit_threshold</code>	Start testing with this parameter set to 50% of the global allocation limit.

Check `M_EXPENSIVE_STATEMENTS.MEMORY_SIZE` for the typical memory usage of statements.

Scenario 4: Large System with at least 4 Sockets - All Workload Types

Large systems usually need to handle many concurrent statements therefore it is usually reasonable to limit the concurrency of statements; but also, this may help to avoid cases where HANA over-parallelizes. Less parallelism per statement on these large systems tends to have better performance because statements run on one single NUMA node and we tend to avoid cross NUMA node communication.

Large System, All Workload Types

Tuning Option	Example Setting
<code>statement_memory_limit</code>	Start testing with this parameter set to 25% of the global allocation limit.
<code>default_statement_concurrency_limit</code>	Assign 33% of the logical cores on the system.
<code>max_concurrency</code>	From SPS12, the default setting should give good performance in most cases. If necessary, consider reducing the value of this parameter. Start testing with the parameter set to 50% of the available logical cores.
<code>max_concurrency_hint, num_cores</code>	Consider setting <code>max_concurrency_hint</code> and <code>num_cores</code> equal to the number of logical cores per socket.
Workload Classes	In mixed scenarios fine-grained control is possible with workload classes.

Related Information

[Understand your Workload \[page 624\]](#)

[Analyzing System Performance \[page 625\]](#)

6.10 Scheduling of Recurring Administration Tasks

SAP HANA Extended Services, classic model (SAP HANA XS classic) provides a job-scheduling feature that allows you to execute an XS JavaScript function or call an SQLScript procedure at a scheduled interval. This may be useful for scheduling recurring administration tasks such as performing backups and running traces at specific times.

You can use the job-scheduling feature of SAP HANA XS classic to run recurring administration tasks in the background at a specified interval. For example:

- Perform backups
- Run traces
- Query specific monitoring views

You create a scheduled job using the `.xsjob` file, a design-time file that you commit to and activate in the SAP HANA repository. The scheduled job can be used to perform the following actions:

- Execute an XS JavaScript function
- Call an SQLScript procedure

Once the job file is available in the required system (for example, after transport from a development system to a production system), you can configure its execution in runtime using the XS Administration Tool. The job-scheduling feature in SAP HANA XS must also be enabled in the system configuration.

To create and enable a recurring task using the job-scheduling feature, you perform the high-level steps listed below. For more detailed information about how to perform the individual steps, see *Scheduling XS Jobs*.

1. Create the function or script that defines the task you want to perform.
2. Create the job file `.xsjob` that defines the details of the recurring task.
3. Maintain the corresponding runtime configuration for the XS job.
4. Enable the job-scheduling feature in SAP HANA XS.
5. Check the job logs to ensure the job is running according to schedule.

Related Information

[Scheduling XS Jobs \[page 1618\]](#)

[Scheduling Jobs in XS Advanced \[page 1790\]](#)

6.11 Getting Support

There are many support resources available to help you work with SAP HANA and a range of tools and functions to help you (together with SAP Support) to analyze, diagnose, and resolve problems.

Support Resources

The SAP Support Portal is the starting point for all support questions; an 'S-user' ID is required for access: <https://support.sap.com>. From here you have access to all resources including the One Support launchpad, support incidents, software downloads, license keys, documentation, SAP Notes and options for contacting Support.

Guided Answers

Guided Answers is an interactive online support tool to help you to diagnose and solve problems using decision trees. It covers many SAP products including SAP HANA and offers a set of step-by-step problem-solving documents each one designed to address a specific issue. Guided Answers is available in the SAP Support portal at the following address: <https://ga.support.sap.com/dtp/viewer/>

i Note

Two Guided Answers trees may be particularly helpful in making an initial assessment of any problem:

- The *SAP HANA Component Guide* tree will help you to troubleshoot an issue and determine the correct component reference (such as *HAN-DB-BAC* for SAP HANA Backup and Recovery). If you need to contact Support the component ID will enable Support analysts to immediately identify where the problem is located. [SAP HANA Component Guide](#)
- The *SAP HANA Troubleshooting and Problem Categorization* tree is a general high-level troubleshooting tree for SAP HANA. [SAP HANA Troubleshooting](#)

SAP Community

In the SAP Community Network (SCN) you can find many support resources online including wikis, blogs, reference materials and so on. This SCN wiki page, for example, provides links to many specialist troubleshooting topics: [SAP HANA In-Memory Troubleshooting Guide](#).

The SAP Questions and Answers knowledge base may also provide quick answers to many common questions: [SAP Q&A](#).

YouTube

Both SAP HANA Academy and SAP Support offer YouTube channels with a wide range of support materials in video format:

- <https://www.youtube.com/user/saphanaacademy>
- <https://www.youtube.com/user/SAPSupportInfo>

Contacting Support

If you need personal interactive support for questions which are not covered by these resources then you may need to open a support incident. The Support team may require you to prepare analysis documents and provide diagnostic information (such as trace files) in order to efficiently analyze your scenario; tools and procedures for doing this are described in detail in this section.

Incidents, Calls and Chat Sessions

In addition to incidents the *Expert Chat* and *Schedule an Expert* support options may be preferable alternatives. A chat session can be either a simple text chat within your browser or you can share screens with an expert to demonstrate a specific issue or question. If you can plan your support call in advance there is also the option of scheduling a 30-minute call with a support engineer. Full details of these options are available in the *My Support* area of the portal: [Incidents](#).

Related Information

[Collecting Diagnosis Information for SAP Support \[page 687\]](#)

[Open a Support Connection \[page 698\]](#)

[SAP Note 2570790](#)

[SAP Note 2392095](#)

[Diagnosis Files \[page 661\]](#)

[Traces \[page 665\]](#)

[Problem Analysis Using hdbcons \[page 698\]](#)

6.11.1 Diagnosis Files

Diagnosis files include log and trace files, as well as a mixture of other diagnosis, error, and information files. In the event of problems with the SAP HANA database, you can check these diagnosis files for errors. Traces generally need to be explicitly enabled and configured.

Location of Diagnosis Files

Diagnosis files of the system database are stored at the following default location on the SAP HANA server: `/usr/sap/<SID>/HDB<instance>/<host>/trace`.

Trace files of tenant databases are stored in a sub-directory named `DB_<database_name>`.

→ Recommendation

Monitor disk space that is used for diagnosis files and delete files that are no longer needed.

Related Information

[Traces \[page 665\]](#)

[Monitoring Disk Space \[page 378\]](#)

[Tenant Databases \[page 19\]](#)

6.11.1.1 View Diagnostic Files in the SAP HANA Database Explorer

Diagnose and analyze errors in an SAP HANA database by viewing the relevant diagnostic files in the SAP HANA database explorer.

Prerequisites

Because the database explorer is integrated with the SAP Web IDE for SAP HANA and the SAP HANA cockpit, you must be a user of one of these applications.

Context

In the database browser, diagnostic files for online databases are grouped by host and then by service. In a multi-host system, check each host folder to view all diagnostic files associated with a particular service.

Procedure

Choose one of the following options:

Option	Description
View diagnostic files for a database that is online and available	Open the <i>Database Diagnostic Files</i> folder of your database, then click the diagnostic file that you want to examine to open it in an editor. The <i>Database Diagnostic Files</i> folder contains all diagnostic files that are available through the database's M_TRACEFILES system view.
View diagnostic files for a cockpit resource that is either online or offline	Open the <i>Host Diagnostic Files</i> folder of your cockpit resource, then click the diagnostic file that you want to examine to open it in an editor. The <i>Host Diagnostic Files</i> folder contains all diagnostic files that have been configured for the SAP Host Agent. For more information about configuring the SAP Host Agent, see the <i>SAP Host Agent</i> documentation. The cockpit resource must have valid SAP Control Credentials set in the cockpit. If the user has not set valid SAP Control Credentials, then an error is returned.
View diagnostic files for tenant databases in an MDC system	Under the <i>Database Diagnostic Files</i> folder of your system, click the folder of the tenant database that you want to view trace files for. Click the diagnostic file that you want to examine to open it in an editor.

i Note

You must be connected to the system database of the MDC system as the SYSTEM user.

6.11.1.2 View Diagnosis Files in SAP HANA Studio

In the event of problems with the SAP HANA database, you can display the relevant diagnosis file(s) in the SAP HANA studio and analyze for errors.

Prerequisites

You have the system privilege CATALOG READ.

Procedure

In the Administration editor, choose the *Diagnosis Files* tab.

All available diagnosis files are displayed. Use the available options to filter the number of files.

i Note

If you are in the system database, you initially see the diagnosis files of the system database only.

For more information about other options for working with diagnosis files, see *Options for Diagnosis File Handling (SAP HANA Studio)*

Related Information

[Options for Diagnosis File Handling \(SAP HANA Studio\) \[page 663\]](#)

[View Diagnosis Files of an Unavailable Tenant Database \[page 264\]](#)

6.11.1.2.1 Options for Diagnosis File Handling (SAP HANA Studio)

In the SAP HANA studio, you can filter, merge, delete, compress, and download diagnosis files.

The following options for working with diagnosis files are available on the *Diagnosis Files* tab of the Administration editor.

Option	Description
Filter files by text	To refine the list of diagnosis files, enter a filter text in the <i>Filter</i> field. For example, if you want to see only SQL trace files, enter sqltrace . You can also filter by host.

Option	Description
Filter files by tenant database (system database only)	To see the diagnosis files for a particular tenant database, select the database name. This filter is only available in the system database. If no tenant database is selected (default), the diagnosis files of the system database are displayed.
Display file contents	<p>To display the contents of a file, right-click it and choose <i>Open</i>, or double-click the file.</p> <p>The <i>Show Start of File</i> and <i>Show End of File</i> buttons help you to navigate particularly large files more easily. You can specify how many lines you want to see when you filter the file in this way.</p> <div data-bbox="603 651 1394 869" style="background-color: #f0f0f0; padding: 5px;"> <p>i Note</p> <p>Depending on the type of data in the diagnosis file, the number of lines actually displayed may be less than or greater than specified. This is because the data in some diagnosis files is fetched in bytes and the number of bytes per line varies.</p> </div> <div data-bbox="603 891 1394 1039" style="background-color: #f0f0f0; padding: 5px;"> <p>→ Tip</p> <p>Crash dump files have a hyperlinked table of contents. To see the hyperlinks, press the <code>CTRL</code> key as you move your mouse over the entries.</p> </div>
Merge file contents	<p>You can merge the diagnosis files listed by choosing <i>Merge Diagnosis Files...</i></p> <p>This feature is helpful during troubleshooting as it allows you to review multiple diagnosis files of different file types at the same time.</p> <p>As merging diagnosis files can take a long time depending on the size and number of files to be merged, you can restrict the amount of data included by specifying a timeframe (for example, data from the last 24 hours only or from between certain dates and times). You can also restrict to data of trace level ERROR.</p> <p>If you select <i>Export to CSV file</i>, the merged trace file is exported to the specified location. If you do not select this option, the file is opened directly in the SAP HANA studio in the <i>Merged Diagnosis Files</i> editor. Here, you can use additional filtering options and the timeframe slider to drill down and analyze further.</p>

Option	Description
Delete files	<p>You delete one or more individual log files or other non-trace files (for example, *.log, *.tpt, *.py) from the list by selecting the file(s) in question and in the context menu, choosing <i>Delete</i>.</p> <p>You delete trace files in the following ways:</p> <ul style="list-style-type: none"> You can delete trace files in the same way as other diagnosis files by right-clicking them and choosing <i>Delete</i>. <div data-bbox="667 584 1390 869" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p>i Note</p> <p>The file may not actually be deleted. If a running service is currently writing to the file, it cannot be deleted. If this is the case, the file disappears from the list in the SAP HANA studio and is hidden in the file system at operating system level. As long as a service is still writing to the file, it still exists and consumes disk space. Once the file reaches its maximum size, the system stops writing to it and creates a new trace file. When the file is physically deleted depends on how trace file rotation is configured.</p> </div> <ul style="list-style-type: none"> You can batch delete trace files, for example all the trace files of a specific service, by choosing <i>Delete Trace Files...</i> and making the required selection. <div data-bbox="667 965 1390 1133" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p>i Note</p> <p>In this case, if it is not possible to delete trace files because they are open, the content of the file is cleared. The file still exists but its size is reduced.</p> </div>
Compress files	<p>If you need to download a diagnosis file (for example, to send it to SAP Support), you can compress it first on the server. This is useful for large diagnosis files and/or slow connections. To compress a file, right-click it and choose <i>Compress</i>. After compression, the file has the file format *.zip. You can select multiple files to compress.</p>
Download files	<p>To download a diagnosis file for offline analysis, right-click and choose <i>Download</i>. You can select multiple files to download.</p>

6.11.2 Traces

SAP HANA provides various traces for obtaining detailed information about the actions of the database system for troubleshooting and error analysis.

The following is an overview of the main traces available in SAP HANA:

Configure or Run...	To...	Using...
Database trace	Understand activity in the components of the SAP HANA database in general, or for a specific user or client operation	SAP HANA database explorer or SAP HANA studio
<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>Database tracing is always active. Information about error situations is always recorded in alert trace files.</p> </div>		
SQL trace	Collect information about all SQL statements executed on the index server	SAP HANA database explorer or SAP HANA studio
Expensive statements	Record information about individual SQL statements whose execution time exceeded a configured threshold	SAP HANA cockpit, SAP HANA database explorer, or SAP HANA studio
Plan trace	Visualize the execution plans of SQL SELECT statements for in-depth query performance analysis	SQL analyzer or SAP HANA PlanViz perspective of the SAP HANA studio
<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>The plan trace is part of the SAP HANA tools for query execution analysis: the Web-based SQL analyzer, which is integrated into the SAP HANA cockpit and SAP Web IDE for SAP HANA, and the SAP HANA PlanViz perspective of the SAP HANA studio.</p> </div>		
Performance trace	Record performance indicators of individual query processing steps in the database kernel for detailed performance analysis by SAP Support	SAP HANA studio
Kernel profiler	Generate a profiler trace for detailed performance analysis by SAP Support	SAP HANA database explorer or SAP HANA studio
Full system information dump	Create a runtime dump if requested by SAP support to understand the behavior of SAP HANA in a problem situation	SAP HANA cockpit or SAP HANA studio
SAP Web Dispatcher HTTP traces	Analyze HTTP requests	SAP HANA cockpit or SAP HANA studio

Related Information

[Database Trace \(Basic, User-Specific, and End-to-End\) \[page 667\]](#)

[SQL Trace \[page 672\]](#)

[Performance Trace \[page 675\]](#)

[Expensive Statements Trace \[page 676\]](#)

[Kernel Profiler \[page 679\]](#)

[Monitor and Analyze Statements with Plan Trace \(SQL Analyzer\) \[page 445\]](#)

[Collecting Diagnosis Information for SAP Support \[page 687\]](#)

[Traces and Trace Configuration for Internal Web Dispatcher \[page 681\]](#)

[Configure Traces in SAP HANA Studio \[page 683\]](#)

[Configure Tracing in the SAP HANA Database Explorer \[page 682\]](#)

6.11.2.1 Database Trace (Basic, User-Specific, and End-to-End)

The database trace records information about activity in the components of the SAP HANA database. You can use this information to analyze performance and to diagnose and debug errors.

Database Trace Files

Each service of the SAP HANA database writes to its own trace file. The file names follow the default naming convention:

```
<service>_<host>.<port_number>.<3_digit_file_counter>.trc.
```

❁ Example

```
indexserver_veadm009.34203.000.trc
```

Information recorded by the `alert` trace component of the database trace is written to a separate file to ensure that critical information is easy to find. The file names follow the default naming convention:

```
<service>_alert_<host>.trc
```

i Note

The trace files generated for the Web Dispatcher service are different. For more information, see *Trace Configuration for Internal Web Dispatcher*.

You access database trace files with other diagnosis files.

Basic Database Trace Configuration

Database tracing is always active. This means that information about error situations is always recorded. However, for more detailed analysis of a specific problem or component, you may need to configure a certain trace component to a trace level higher than the default value. For example, the `backup` trace component records information about backup operations, the `authorization` component records information about authorization operations, and so on.

→ Tip

For more information about which trace component to use for which situation, see SAP Note 2380176.

Changing the trace level of a trace component

If a trace component is available in all services, the trace level can be configured for all services at once. It is also possible to configure the trace level of a component individually for a specific service. The trace level of a component configured at service level overrides the trace level configured for all services. Some components are only available in a particular service and cannot therefore be changed globally.

❁ Example

You change the trace level of the `memory` component to `ERROR` for all services and for the `indexserver` service, you change it to `WARNING`. This means that the `memory` component of the `indexserver` service will trace up to level `WARNING` and the `memory` component of all other services will trace to the level `ERROR`.

You can configure the trace levels of database trace components in the SAP HANA database explorer or the SAP HANA studio. Alternatively, you can modify the parameters in the `trace` section of the `global.ini` configuration file (for all services) or service-specific files such as `indexserver.ini`. The individual parameters correspond to trace components and the parameter value is the trace level.

❁ Example

Use the following statement to set the trace level of the `authentication` component to `DEBUG`:

```
ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini', 'SYSTEM') SET ('trace',  
'authentication') = 'DEBUG' WITH RECONFIGURE
```

i Note

In the SAP HANA database explorer and SAP HANA studio, a trace level that has been inherited from the global or *ALL SERVICES* configuration (either the default or system configuration) is shown in brackets.

What trace levels are possible?

The higher the trace level, the more detailed the information recorded by the trace. The following trace levels exist:

- NONE (0)
- FATAL (1)
- ERROR (2)
- WARNING (3)
- INFO (4)
- DEBUG (5)

i Note

Even if you select trace level `NONE`, information about error situations is still recorded.

Alert Trace Component

The `alert` trace component is used to ensure that critical information is easy to find by duplicating high-priority trace messages to separate alert trace files. By default, trace messages with trace level ERROR and higher are written to the alert trace file.

❖ Example

If the `alert` component is set to WARNING, the alert trace contains messages created with trace levels WARNING, FATAL and ERROR.

User-Specific and End-to-End Database Traces

User-specific and end-to-end traces extend the configured database trace by allowing you to change the trace level of components in the context of a particular user or end-to-end analysis. The trace levels configured for components in these contexts override those configured in the database trace.

❖ Example

In the database trace, you changed the trace level of the `memory` component to ERROR for all services, and for the `indexserver` service you changed it to WARNING. Now, you create a user-specific trace for User1 and increase the trace level for all services to WARNING. For the `indexserver` service, you increase it to DEBUG. This results in the following tracing behavior for the `memory` component:

- For all users except User1, all services except the `indexserver` will trace to ERROR
- For all users except User1, the `indexserver` will trace to WARNING
- For User1, all services except the `indexserver` will trace to WARNING
- For User1, the `indexserver` will trace to DEBUG

End-to-End Traces

End-to-end traces are triggered by applications outside of the SAP HANA database. The default trace levels for the SAP HANA database components are normally sufficient and do not need to be changed. For more information about end-to-end analysis in your landscape, see SAP Library for Solution Manager on SAP Help Portal.

Related Information

[Diagnosis Files \[page 661\]](#)

[Traces and Trace Configuration for Internal Web Dispatcher \[page 681\]](#)

[End-to-End Analysis Overview](#)

[SAP Note 2380176](#)

6.11.2.1.1 Database Trace Configuration in Tenant Databases

Tenant databases inherit the database trace level configured in the system database unless you change the trace level in the tenant database.

The trace level of trace components in a tenant database is inherited from the system database as the default value. If you want to configure a different trace level for a particular component in the tenant database, either globally for all services or for a specific service, you can do so by changing the trace level of the relevant component.

i Note

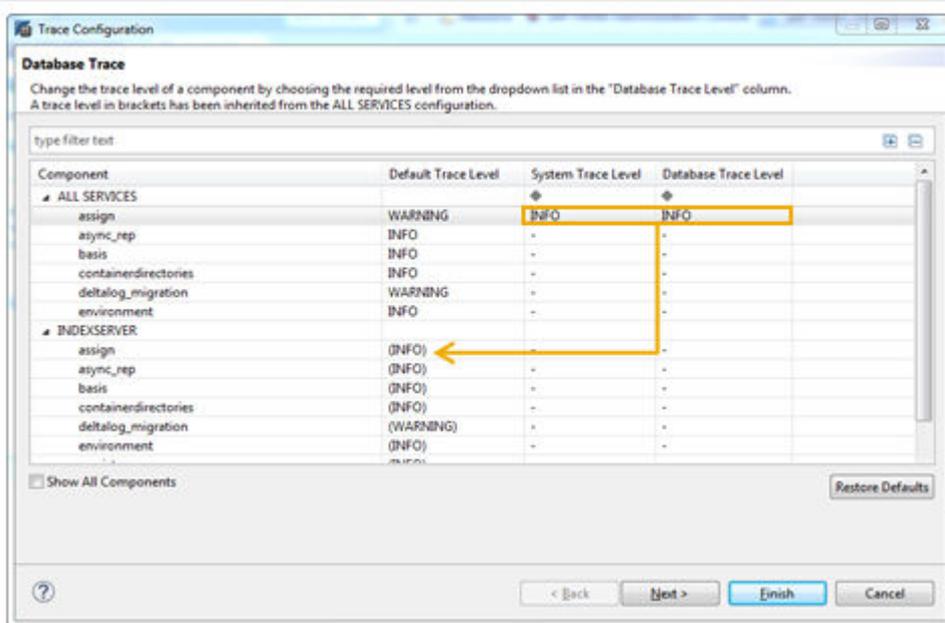
In the SAP HANA studio, the trace level of a component will also be displayed as the system trace level, and you cannot change the system trace level. This is because from the perspective of the tenant database, the database and the system are effectively the same. The true system trace level (that is, the value configured in the system database) appears as the default trace level for the tenant database.

i Note

A trace level that has been inherited from the ALL SERVICES configuration (either the default or database configuration) is shown in brackets.

❖ Example

In the following example, you can see that the default trace level for the component `assign` inherited from the system database is `WARNING`. The trace level was changed to `INFO` for all services in this database. The `indexserver` service therefore inherits this trace level. All the other components keep the default configuration.



Database Trace Configuration in a Tenant Database (Studio View)

6.11.2.1.2 Configure Database Trace File Rotation

Trace file rotation prevents trace files from growing indefinitely by limiting the size and number of trace files. You can configure trace file rotation globally for all services in the database and for individual services.

Prerequisites

You have the system privilege INIFILE ADMIN.

Context

You configure trace file rotation in the `global.ini` configuration file (all services) or the service-specific file (for example, `indexserver.ini`).

Procedure

Depending on whether you are configuring trace file rotation for all system services or for an individual service, proceed as follows:

Option	Description
All services	<p>In the <code>trace</code> section of the <code>global.ini</code> file, configure the following parameters:</p> <ul style="list-style-type: none">◦ <code>maxfiles</code> by specifying the maximum number of trace files that may exist◦ <code>maxfilesize</code> by specifying in bytes the maximum size an individual trace file may reach <div data-bbox="331 1507 1394 1641"><p>i Note</p><p>The default configuration for trace file rotation in the <code>global.ini</code> file is <code>maxfiles=10</code> and <code>maxfilesize=10000000</code>.</p></div>
Individual service	<p>In the <code>trace</code> section of the <code>global.ini</code> file, configure the following parameters. If there is no <code>trace</code> section, add one.</p> <ul style="list-style-type: none">◦ <code>maxfiles</code> by specifying the maximum number of trace files that may exist◦ <code>maxfilesize</code> by specifying in bytes the maximum size an individual trace file may reach in bytes <div data-bbox="331 1805 1394 1939"><p>i Note</p><p>If these two parameters do not exist in the <code>trace</code> section or if you created a new <code>trace</code> section, create them.</p></div>

Results

When a trace file reaches the specified maximum file size, it is closed, and a new file created. When the specified maximum number of files is reached, the next time a new file is created, the first file is deleted, and so on.

i Note

The system checks the size and number of diagnosis files regularly. The threshold values for these checks (check 50 and 51) should be in line with the configured trace file rotation.

Related Information

[Configuring SAP HANA System Properties \(INI Files\) \[page 291\]](#)

[Configure Traces in SAP HANA Studio \[page 683\]](#)

[Configure Check Thresholds \[page 377\]](#)

6.11.2.2 SQL Trace

The SQL trace collects information about all SQL statements executed on the index server (tenant database) or name server (system database) and saves it in a trace file for further analysis. The SQL trace is inactive by default.

Information collected by the SQL trace includes overall execution time of each statement, the number of records affected, potential errors (for example, unique constraint violations) that were reported, the database connection being used, and so on. The SQL trace is a good starting point for understanding executed statements and their potential effect on the overall application and system performance, as well as for identifying potential performance bottlenecks at statement level.

SQL Trace Files

SQL trace information is saved as an executable python program (by default `sqltrace_<...>.py`), which can be accessed with other diagnosis files.

Enabling and Configuring the SQL Trace

You can enable and configure the SQL trace in the SAP HANA database explorer or SAP HANA studio. Alternatively, you can modify the parameters in the `sqltrace` section of the `indexserver.ini` (tenant database) or `nameserver.ini` (system database).

❁ Example

Use the following statement to enable the SQL trace:

```
ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini', 'SYSTEM') SET ('sqltrace',  
'trace') = 'on' WITH RECONFIGURE
```

→ Recommendation

Do not leave the SQL trace enabled all the time as writing trace files consumes storage space on the disk and can impact database performance significantly.

Trace Levels

You can configure the following trace levels for the SQL trace. The trace level corresponds to the configuration parameter [sqltrace] level in the indexserver.ini file (tenant database) or nameserver.ini file (system database).

Trace Level	Description
NORMAL	All statements that have finished successfully are traced with detailed information such as executed timestamp, thread ID, connection ID, and statement ID.
ERROR	All statements that returned errors are traced with detailed information such as executed timestamp, thread ID, connection ID, and statement ID.
ERROR_ROLLBACK	All statements that are rolled back are traced with detailed information such as executed timestamp, thread ID, connection ID and statement ID.
ALL	All statements including status of normal, error, and rollback are traced with detailed information such as executed timestamp, thread ID, connection ID and statement ID.
ALL_WITH_RESULTS	In addition to the trace generated with trace level ALL, the result returned by select statements is also included in the trace file.

i Note

An SQL trace that includes results can quickly become very large.

⚠ Caution

An SQL trace that includes results may expose security-relevant data, for example, query result sets.

Additional Configuration Options

Option	Configuration Parameter	Default	Description
Trace file name	<code>tracefile</code>	<code>sqltrace_</code> <code>\$HOST_</code> <code>{PORT}_</code> <code>{COUNT:</code> <code>3}.py</code>	<p>User-specific name for the trace file</p> <p>If you do not enter a user-specific file name, the file name is generated according to the following default pattern:</p> <p><code>DB_<dbname>/sqltrace_</code> <code>\$HOST_</code> <code>{PORT}_</code> <code>{COUNT:3}.py</code>, where:</p> <ul style="list-style-type: none"> • <code>DB_<dbname></code> is the sub-directory where the trace file is written if you are running on a tenant database • <code>\$HOST</code> is the host name of the service (for example, <code>indexserver</code>) • <code>\$PORT</code> is the port number of the service • <code>\$COUNT:3</code> is an automatically generated 3-digit number starting with 000 that increments by 1 and serves as a file counter when several files are created.
User, application, object, and statement filters	<code>user</code> <hr/> <code>application_user</code> <hr/> <code>application</code> <hr/> <code>object</code> <hr/> <code>statement_type</code>	Empty string	<p>Filters to restrict traced statements to those of particular database or application users and applications, as well as to certain statement types and specific objects (tables, views, procedures).</p> <p>All statements matching the filter criteria are recorded and saved to the specified trace file.</p>
Flush limit	<code>flush_interval</code>	16	<p>During tracing, the messages of a connection are buffered. As soon as the flush limit number of messages is buffered (or if the connection is closed), those messages are written to the trace file.</p> <p>When set to 0, every SQL trace statement is immediately written to the trace file</p>

Trace File Rotation

The size and number of trace files are controlled by the following parameters.

Parameter	Default	Description
<code>max_files</code>	1	Sets the maximum number of trace files
<code>filesize_limit</code>	1610612736 (or 1.5 GB)	Sets the maximum size of an individual trace file in bytes

⚠ Caution

If both the maximum number of files and the maximum file size are reached, SQL tracing stops. If this happens, you can increase the values of `max_files` and `filesize_limit`. See SAP Note 2629103.

SAP HANA SQL Trace Analyzer

SAP HANA SQL trace analyzer is a Python tool you can use to analyze the HANA SQL trace output. The tool gives you an overview of the top SQL statements, the tables accessed, statistical information on different statement types and on transactions executed.

For more information about the installation and usage of SAP HANA SQL trace analyzer, see SAP Knowledge Base Article 2412519 *FAQ: SAP HANA SQL Trace Analyzer* .

Related Information

[Diagnosis Files \[page 661\]](#)

[SAP Note 2036111](#) 

[SAP Note 2412519](#) 

[SAP Note 2629103](#) 

6.11.2.3 Performance Trace

The performance trace is a performance tracing tool built into the SAP HANA database. It records performance indicators for individual query processing steps in the database kernel. You may be requested by SAP Support to provide a performance trace.

Information collected includes the processing time required in a particular step, the data size read and written, network communication, and information specific to the operator or processing-step-specific (for example, number of records used as input and output). The performance trace can be enabled in multiple tenant databases at the same time to analyze cross-database queries.

Performance Trace Files

Performance trace results are saved to the trace files with file extension *.tpt or *.cpt, which you can access with other diagnosis files. To analyze these files, you need a tool capable of reading the output format (*.tpt and *.cpt). SAP Support has tools for evaluating performance traces.

Enabling and Configuring the Performance Trace

You can enable and configure the performance trace in the SAP HANA studio or using the `ALTER SYSTEM * PERFTRACE` SQL statements.

Example

To start the performance trace execute `ALTER SYSTEM START PERFTRACE`.

Configuration Options

Option	Description
Trace file name	The name of the file to which the trace data is automatically saved after the performance trace is stopped
User and application filters	Filters to restrict the trace to a single specific database user, a single specific application user, and a single specific application
Trace execution plans	You can trace execution plans in addition to the default trace data.
Function profiler	The function profiler is a very fine-grained performance tracing tool based on source code instrumentation. It complements the performance trace by providing even more detailed information about the individual processing steps that are done in the database kernel.
Duration	How long you want tracing to run If a certain scenario is to be traced, ensure that you enter a value greater than the time it takes the scenario to run. If there is no specific scenario to trace but instead general system performance, then enter a reasonable value. After the specified duration, the trace stops automatically.

Additional filter options are available in extended mode to restrict the trace data further.

For more information about how to configure the performance trace using SQL, see the *SAP HANA SQL and System Views Reference*.

6.11.2.4 Expensive Statements Trace

Expensive statements are individual SQL statements whose execution time exceeds a configured threshold. The expensive statements trace records information about these statements for further analysis and is inactive by default.

If, in addition to activating the expensive statements trace, you enable per-statement memory tracking, the expensive statements trace will also show the peak memory size used to execute the expensive statements.

Expensive Statements Trace Information

If you have the TRACE ADMIN privilege, then you can view expensive statements trace information in the following ways:

- In the *Expensive Statements* app of the SAP HANA cockpit
- On the **Trace Configuration** > *Expensive Statements Trace* tab of the SAP HANA database explorer
- On the **Performance** > *Expensive Statements* tab of the SAP HANA studio
- In the M_EXPENSIVE_STATEMENTS system view

Enabling and Configuring the Expensive Statements Trace

You can enable and activate the expensive statements trace in the SAP HANA cockpit or the SAP HANA database explorer. Alternatively, you can modify the parameters in the `expensive_statement` section of the `global.ini` configuration file.

Configuration Options

i Note

The following table shows the configuration parameters which are available; not all of these are available in the SAP HANA cockpit or the SAP HANA database explorer.

Option	Configuration Parameter	Default Value	Description
Trace status	<code>enable</code>	off	Specifies the activation status of the trace.
Threshold CPU time	<code>threshold_cpu_time</code>	-1 (disabled)	Specifies the threshold CPU time of statement execution in microseconds. When set to 0, all SQL statements are traced.
Threshold memory	<code>threshold_memory</code>	-1 (disabled)	Specifies the threshold memory usage of statement execution in bytes. When set to 0, all SQL statements are traced.
Threshold duration	<code>threshold_duration</code>	1000000 (microseconds = 1 second)	Specifies the threshold execution time in microseconds. When set to 0, all SQL statements are traced. In the SAP HANA database explorer, you can set the threshold duration to be measured in seconds or milliseconds.
User, application, and object filters	<code>user</code> <code>application_user</code>	Empty string	Specifies filters to restrict traced statements to those of a particular database, application user, application, or tables/views.

i Note

Resource tracking and CPU time tracking must also be enabled. You can do this by configuring the corresponding parameters in the `resource_tracking` section of the `global.ini` file.

i Note

Resource tracking and memory tracking must also be enabled. You can do this by configuring the corresponding parameters in the `resource_tracking` section of the `global.ini` file.

Option	Configuration Parameter	Default Value	Description
	application		
	object		
Passport trace level	passport_tracelevel	Empty string	<p>If you are activating the expensive statements trace as part of an end-to-end trace scenario with the Process Monitoring Infrastructure (PMI), you can specify the passport trace level as an additional filter.</p> <p>This means that only requests that are marked with a passport of the specified level are traced.</p>
<div style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 5px;"> <p>i Note</p> <p>Process tracing is possible only for components in the ABAP and Business Objects stacks.</p> </div>			
Trace parameter values	trace_parameter_values	true	In SQL statements, field values may be specified as parameters (using a "?" in the syntax). If these parameter values are not required, then you can disable this setting to reduce the amount of data traced.
Trace flush interval	trace_flush_interval	10	Specifies the number of records after which a trace file is flushed.
Use in-memory tracing	use_in_memory_tracing	true	If in-memory tracing is active, then information is cached in memory. Otherwise, the data is written directly to file.
In-memory tracing records	in_memory_tracing_reco rds	30000	<p>Specifies the maximum number of trace records (per service) stored in memory.</p> <p>This setting only takes effect when in memory tracing is active.</p>

Trace File Rotation

To prevent expensive statement trace information from growing indefinitely, you can limit the size and number of trace files using the following parameters in `expensive_statement` of `global.ini`.

Parameter	Default	Description
maxfiles	10	<p>Specifies the maximum number of trace files.</p> <p>When the maximum number of trace files reached, the oldest trace file is deleted and a new one opened.</p> <p>When set to 0, trace file rotation is disabled.</p>

Parameter	Default	Description
maxfilesize	10000000 (or 9.5 megabytes)	<p>Specifies the maximum size of an individual trace file in bytes.</p> <p>When the maximum number of files is greater than 1 and the maximum file size is reached, a new trace file is opened.</p> <p>When the maximum number of files is 1, the maximum file size is greater than zero, and the maximum file size is reached, the trace file is deleted and a new one opened.</p>

Related Information

[Setting a Memory Limit for SQL Statements \[page 633\]](#)

[Monitoring and Analyzing Expensive Statements \(SAP HANA Cockpit\) \[page 407\]](#)

[Expensive Statements Monitoring \(SAP HANA Studio\) \[page 456\]](#)

[SAP Note 2180165](#) 

[SAP Note 2036111](#) 

6.11.2.5 Kernel Profiler

The kernel profiler is a sampling profiler built into the SAP HANA database. It can be used to analyze performance issues with systems on which third-party software cannot be installed, or parts of the database that are not accessible by the performance trace. It is inactive by default.

The kernel profile collects, for example, information about frequent and/or expensive execution paths during query processing.

Kernel Profiler Traces

Profiling results are saved to the trace files `CPU_<service>_<host>_<port>_<timestamp>.<format>` and `WAIT_<service>_<host>_<port>_<timestamp>.<format>`, where `<format>` is either dot or `kcachegrind`. You can access the profiling results with other diagnosis files. To analyze these trace files meaningfully you need a tool capable of reading the configured output format, that is `KCacheGrind` or `DOT` (default format). Alternatively, send the files to SAP Support.

Enabling and Configuring the Kernel Profiler

You enable and configure the kernel profile in the SAP HANA studio or the SAP HANA database explorer. It is recommended that you start kernel profiler tracing immediately before you execute the statements you want to analyze and stop it immediately after they have finished. This avoids the unnecessary recording of irrelevant statements. It is also advisable as this kind of tracing can negatively impact performance.

i Note

To enable the kernel profile, you must have the SAP_INTERNAL_HANA_SUPPORT role. This role is intended only for SAP HANA development support.

Configuration Options

Option	Description
Service(s) to profile	The service(s) that you want to profile
Wait time	The amount of time the kernel profiler is to wait between call stack retrievals When you activate the kernel profiler, it retrieves the call stacks of relevant threads several times. If a wait time is specified, it must wait the specified time minus the time the previous retrieval took.
Memory limit	Memory limit that will stop tracing The kernel profiler can potentially use a lot a memory. To prevent the SAP HANA database from running out of memory due to profiling, you can specify a memory limit that cannot be exceeded.
Database user, application user	The database user or application user you want to profile
Use KCachegrind format to write trace files	Output format of trace files (configurable when you stop tracing)

Related Information

[Diagnosis Files \[page 661\]](#)

6.11.2.6 Traces and Trace Configuration for Internal Web Dispatcher

Several traces and trace configuration options are available for the internal Web Dispatcher, which runs as a native SAP HANA service (`webdispatcher`).

Developer Trace

The developer trace is the main trace for the Web Dispatcher and contains technical information for troubleshooting problems.

The developer trace file is `webdispatcher_<host>.<port>_dev_webdisp`.

You can configure the developer trace in the following ways:

- Changing the database trace level for the `dev_webdisp` component of the `webdispatcher` service
The default trace level is `ERROR`.
- Changing (or adding) the property `rdisp/trace` in the `[profile]` section of the `webdispatcher.ini` configuration file.
Possible values are 0, 1, 2, and 3.

Database Trace

The database trace files for the Web Dispatcher contain secondary information related to the Web Dispatcher's integration into the SAP HANA integration system (start/stop, configuration changes, and so on).

The database trace files are:

- `webdispatcher_<host>.<port>.<3_digit_file_counter>.trc`
- `webdispatcher_alert.<host>.trc`

You can configure the database trace by changing the trace level for the `webdispatcher` component of the `webdispatcher` service.

Header Trace

The header trace allows you to analyze HTTP requests and responses efficiently since it contains only the request data and no information about the internal workings of Web Dispatcher.

You can activate the header trace by adding the property `icm/http/trace_info` in the `[profile]` section of the `webdispatcher.ini` configuration file and setting the value to `true`. The trace level is `false` by default.

Header trace information is written to the `dev_webdisp` trace file.

HTTP Access Log

To monitor all HTTP(s) requests processed in an SAP HANA system, you can set up the internal Web Dispatcher to write a standardized HTTP log for each request.

To configure the Web Dispatcher to log all HTTP(s) requests, you add the property `icm/http/logging_0` to the `[profile]` section of the `webdispatcher.ini` configuration file, specifying the following value:

```
PREFIX=/, LOGFILE=$(DIR_INSTANCE)/trace/access_log-%y-%m-%d, MAXSIZEKB=10000,  
SWITCHTF=day, LOGFORMAT=SAP
```

The access log file is `access_log-<timestamp>`.

❁ Example

```
Sample log file entry: [26/Nov/2014:13:42:04 +0200] 10.18.209.126 BOB - "GET /sap/xse/  
test/InsertComment.xsjs HTTP/1.1" 200 5 245
```

The last three numbers are the HTTP response code, the response time in milliseconds, and the size in bytes. For more information about logging and alternative log formats, see the Internet Communication Manager (ICM) documentation on SAP Help Portal.

Related Information

[Database Trace \(Basic, User-Specific, and End-to-End\) \[page 667\]](#)

[rdisp/TRACE* Parameters](#)

[icm/HTTP/trace_info](#)

[icm/HTTP/logging_<xx>](#)

[Logging in the ICM and SAP Web Dispatcher](#)

6.11.2.7 Configure Tracing in the SAP HANA Database Explorer

Various types of tracing are available in the SAP HANA database explorer for obtaining detailed information about the actions of your database system.

Prerequisites

To configure tracing, you must have the TRACE ADMIN system privilege.

Procedure

1. Right-click your database and click *Trace Configuration*.
2. Click *Edit* on the type of tracing that you want to configure.
3. On the *Trace Configuration* editor, configure your trace settings and click *OK*.

Results

You have configured tracing and can view and download tracing results from either your *Database Diagnostic Files* folder or your *Host Diagnostic Files* folder.

6.11.2.8 Configure Traces in SAP HANA Studio

Various traces are available for obtaining detailed information about the actions of the database system. You can activate and configure traces in the SAP HANA studio on the *Trace Configuration* tab of the Administration editor.

Prerequisites

To configure traces, you must have the system privilege TRACE ADMIN. To configure the kernel profiler, you must have the SAP_INTERNAL_HANA_SUPPORT standard role.

Context

You can configure the following traces:

- Database traces (basic, user, end-to-end)
- SQL trace
- Performance trace
- Expensive statements trace
- Kernel profiler
- Plan trace

Procedure

1. In the Administration editor, choose the *Trace Configuration* tab.

2. Choose the  (*Edit Configuration*) button for the trace that you want to configure. The *Trace Configuration* dialog box appears.
3. Make the required settings.

The configuration options available in the *Trace Configuration* dialog box depend on the trace type.

i Note

To restore the default status or configuration of a trace, in the *Trace Configuration* dialog box choose *Restore Defaults*.

→ Tip

If you are configuring the database trace, not all trace components are visible by default in the *Trace Configuration* dialog box. To view all additional components, select *Show All Components*.

Results

Data recorded by traces is saved to trace files, which you can view on the *Diagnosis Files* tab.

Related Information

[Database Trace \(Basic, User-Specific, and End-to-End\) \[page 667\]](#)

[SQL Trace \[page 672\]](#)

[Performance Trace \[page 675\]](#)

[Expensive Statements Trace \[page 676\]](#)

[Kernel Profiler \[page 679\]](#)

6.11.3 Troubleshooting an Inaccessible or Unresponsive SAP HANA System

For situations when a system cannot be reached by SQL or is experiencing performance problems, both the SAP HANA studio and the SAP HANA cockpit provide mechanisms by which you or an SAP support engineer can access diagnosis information and perform emergency operations to resolve the situation.

Related Information

[Using the Cockpit to Troubleshoot an Unresponsive Resource \[page 685\]](#)

[Troubleshoot Unresponsive System in SAP HANA Studio \[page 685\]](#)

6.11.3.1 Using the Cockpit to Troubleshoot an Unresponsive Resource

Using SAP HANA cockpit, you can identify performance issues in a specific system or a tenant database by examining transactional information.

Once you have opened SAP HANA cockpit, you can open *Troubleshoot Unresponsive System* by clicking the corresponding *Administration* link in the system *Overview*.

→ Tip

From *My Resources*, which is displayed when you first open the cockpit, selecting a resource name in the *Recently Accessed* list allows you to drill down to the *Overview* for that individual resource.

Troubleshoot Unresponsive System organizes information about the system by tab. You can diagnose:

- Connections
- Transactions
- Blocked Transactions
- Threads

You can select one connection or transaction to cancel, or choose to cancel all of them.

i Note

To access this page and trigger the collection of the above information, you require the credentials of the `<sid>adm` user. These can be entered for the system or tenant database in the *Resource Directory*.

Related Information

[Open SAP HANA Cockpit \[page 47\]](#)

[Using the Overview to Manage a Resource \[page 318\]](#)

[Working with the Resource Directory \[page 165\]](#)

6.11.3.2 Troubleshoot Unresponsive System in SAP HANA Studio

When a system cannot be reached by SQL or is experiencing performance problems, you cannot use the Administration editor to troubleshoot and/or resolve issues. By opening the Administration editor in diagnosis

mode you or an SAP support engineer can access diagnosis information and perform emergency operations to resolve the situation.

Prerequisites

You have the credentials of the operating system user (<sid>adm user) that was created when the system was installed.

Procedure

1. Open the Administration editor in diagnosis mode:
 - If the system is stopped or cannot be reached by SQL, double-click the system in the *Systems* view.
 - If the system is running, choose the  (*Open Diagnosis Mode*) button from the drop-down menu of the  (*Administration*) button in the *Systems* view.

i Note

The Administration editor is available in diagnosis mode only from the system database. If a tenant database is unavailable, you can view its diagnosis files in the standard Administrator editor of the system database on the *Diagnosis Files* tab.

2. If required, enter the <sid>adm user name and password.

Results

The Administration editor opens in diagnosis mode.

Here, you see the operational status of all services in the system (*Processes* tab) and you have access to log and trace files (*Diagnosis Files* tab). It is also possible to trigger the collection of diagnosis information into a zip file, which you can then download and attach to a support message. For more information, see *Collect and Download Diagnosis Information in SAP HANA Studio*.

If transactional problems are the source of performance issues, you can analyze current activity in the system on the *Emergency Information* tab. Here you see all connections, transactions, blocked transactions, and threads in the system. If necessary, you can cancel individual connections and transactions, or even cancel all transactions.

Related Information

[Collect and Download Diagnosis Information in SAP HANA Studio \[page 689\]](#)

6.11.4 Collecting Diagnosis Information for SAP Support

To help SAP Support analyze and diagnose problems with your system, you can collect a range of diagnosis information from your system into a zip file. You can trigger the collection of diagnosis information from the SAP HANA cockpit, the SAP HANA studio, and the command line.

Related Information

[Collect and Download Diagnosis Information with the Cockpit \[page 687\]](#)

[Collect and Download Diagnosis Information in SAP HANA Studio \[page 689\]](#)

[Collect Diagnosis Information from the Command Line \[page 691\]](#)

[fullSystemInfoDump.py Command Line Options \[page 692\]](#)

[Diagnosis Information Collected \[page 694\]](#)

6.11.4.1 Collect and Download Diagnosis Information with the Cockpit

To help SAP Support analyze and diagnose problems with the SAP HANA database, you can collect diagnosis information into a zip file, which you can then download and attach to a support message for example. With the SAP HANA cockpit, you can create and manage system information dumps.

Prerequisites

- If the database is online, you need the following privileges:

To...	You Need...
Collect diagnosis information	EXECUTE privilege on the procedure <code>SYS.FULL_SYSTEM_INFO_DUMP_CREATE</code>
List diagnosis information	SELECT privilege on the view <code>SYS.FULL_SYSTEM_INFO_DUMPS</code> In the system database of a multiple-container system, you also need SELECT on <code>SYS_DATABASES.FULL_SYSTEM_INFO_DUMPS</code> so that you can see diagnosis information collected from tenant databases.
Download collected diagnosis information	EXECUTE privilege on the procedure <code>SYS.FULL_SYSTEM_INFO_DUMP_RETRIEVE</code>

To...	You Need...
Delete collected diagnosis information	EXECUTE privilege on the procedure SYS.FULL_SYSTEM_INFO_DUMP_DELETE

- If the system is online, but you want to switch it to offline before collecting information, you will be prompted to connect to the resource using the SAP Control credentials.
- If the system is offline (including the system database in a multiple-container system), you must have credentials of the operating system administrator (user <sid>adm).
- If the database is a tenant database in a multiple-container system and it is offline, you must be logged on to the system database and have the privileges listed above. It is not possible to collect, list, download, or delete diagnosis information from an offline tenant database.

Procedure

1. On the system overview page, under *Alerting and Diagnostics*, select *Manage full system information dumps*.
2. On the *Diagnosis Files* page, if the system is online, you can use the drop down list to switch to offline. You will be prompted to connect to the resource with the SAP Control credentials. (If the system is offline, you cannot switch to online).
3. On the *Diagnosis Files* page, choose a zip file from the list or click *Collect Diagnostics* to create a new zip file.
4. When creating a new zip file, specify the scope of information to be collected:

Option	Description
Collect from existing files	Select this option if you want to collect diagnosis information for one or more file types, for a specific time period, by default the last 7 days. If you also want information from system views, then select <i>Include system views</i> .
	<p>i Note</p> <p>If you are connected to the system database of a multiple-container system, only information from the system views of the system database will be collected. Information from the system views of tenant databases will not be collected regardless of this option.</p> <p>Information from system views is collected through the execution of SQL statements, which may impact performance. In addition, the database must be online, so this option is not available in diagnosis mode.</p>
Create and collect one or multiple sets of runtime environment (RTE) dump files	Select this option if you want to restrict the information collected to one or more RTE dump files. You can configure the creation and collection of dump files by specifying the following additional information: <ul style="list-style-type: none"> ○ The number sets to be collected (that is, the number of points in time at which RTE dump files will be collected). Possible values are 1- 5. ○ The interval (in minutes) at which RTE dump files are to be collected (possible values are 1, 5, 10, 15, and 30). The default value is 1. ○ The host(s) from which RTE dump files are to be collected. ○ The service(s) for each selected host from which RTE dump files are to be collected. ○ The section(s) from each selected service from which RTE dump files are to be collected.

The system collects the relevant information and saves it to a zip file. This may take some time and can be allowed to run in the background.

If you are connected to the system database of a multiple-container system, information from all tenant databases is collected and saved to separate zip files.

6.11.4.2 Collect and Download Diagnosis Information in SAP HANA Studio

To help SAP Support analyze and diagnose problems with the SAP HANA database, you can collect diagnosis information into a zip file, which you can then download and attach to a support message for example. The SAP HANA studio uses either SQL or the SAP Host Agent to collect diagnosis information depending on whether SAP HANA is either online or offline.

Prerequisites

- If the database is online, you need the following privileges:

To...	You Need...
Collect diagnosis information	EXECUTE privilege on the procedure <code>SYS.FULL_SYSTEM_INFO_DUMP_CREATE</code>
List diagnosis information	SELECT privilege on the view <code>SYS.FULL_SYSTEM_INFO_DUMPS</code> In the system database of a multiple-container system, you also need SELECT on <code>SYS_DATABASES.FULL_SYSTEM_INFO_DUMPS</code> so that you can see diagnosis information collected from tenant databases.
Download collected diagnosis information	EXECUTE privilege on the procedure <code>SYS.FULL_SYSTEM_INFO_DUMP_RETRIEVE</code>
Delete collected diagnosis information	EXECUTE privilege on the procedure <code>SYS.FULL_SYSTEM_INFO_DUMP_DELETE</code>

- If the database is a tenant database in a multiple-container system and it is currently offline, you must be logged on to the system database and have the privileges listed above. It is not possible to collect, list, download, or delete diagnosis information from an offline tenant database.
- If the system is offline (including the system database in a multiple-container system), you must have credentials of the operating system administrator (user `<sid>adm`). It is not possible to collect, list, download, or delete diagnosis information via an SQL connection.

Procedure

1. In the Administration editor, choose the *Diagnosis Files* tab.

i Note

If there is no connection to the database, the Administration editor opens in diagnosis mode and you will be prompted to enter the credentials of the <sid>adm user. If you are a tenant database administrator and there is no connection to your tenant database, you cannot proceed. Only the system administrator can collect diagnosis information from the system database.

2. Choose ► *Diagnosis Information* ► *Collect* ▾.
3. Specify the scope of information to be collected:

Option	Description
Collect all diagnosis information	Select this option if you want to collect all diagnosis information for a specific time period, by default the last 7 days. If you also want information from system views, then select <i>Include system views</i> . i Note If you are connected to the system database of a multiple-container system, only information from the system views of the system database will be collected. Information from the system views of tenant databases will not be collected regardless of this option. Information from system views is collected through the execution of SQL statements, which may impact performance. In addition, the database must be online, so this option is not available in diagnosis mode.
Create and collect one or multiple sets of runtime environment (RTE) dump files	Select this option if you want to restrict the information collected to one or more RTE dump files. You can configure the creation and collection of dump files by specifying the following additional information: <ul style="list-style-type: none">○ The index server(s) from which RTE dump files are to be collected○ The number of RTE dump file sets to be collected (possible values are 1, 2, 3, 4, and 5)○ The interval (in minutes) at which RTE dump files are to be collected (possible values are 1, 5, 10, 15, and 30). The default value is 1.

i Note

Older systems do not support all of the above options. It may not be possible to exclude system views from collection or you may require operating system (<sid>adm) user access to do so.

The system collects the relevant information and saves it to a zip file. This may take some time and can be allowed to run in the background.

If you are connected to the system database of a multiple-container system, information from all tenant databases is collected and saved to separate zip files.

4. To download the zip file containing the collected diagnosis information, proceed as follows:
 - a. Choose ► *Diagnosis Files* ► *List* ▾.
The *Diagnosis Information* dialog box opens. The zip file containing the collected diagnosis information is listed together with any other zip files of previously collected information.
 - b. Select the relevant zip file and choose *Download Collection*.
 - c. Specify the download location.

- Optional: Delete any old collections that you no longer need by selecting them and choosing [Delete Collections](#).

Related Information

[Diagnosis Information Collected \[page 694\]](#)

6.11.4.3 Collect Diagnosis Information from the Command Line

The `fullSystemInfoDump.py` script allows you to collect information from your system, even when it is not accessible by SQL. You can then add this information to a support message, for example. The script is part of the SAP HANA server installation and can be executed directly from the command line.

Prerequisites

You are logged on as the operating system user, `<sid>adm`.

Context

The `fullSystemInfoDump.py` script is part of the server installation and can be run from the command line. It is located in the directory `$DIR_INSTANCE/exe/python_support`.

i Note

In a multiple-container system, only the system administrator can collect diagnosis information from the command line since tenant database administrators do not have operating system access. Tenant database administrators must use the SAP HANA cockpit or studio to collect diagnosis information from their database.

Procedure

Start the script from its location with the command:

```
python fullSystemInfoDump.py
```

You can modify the command with several command line options. To see the available options, specify the option `--help`.

If the system can be reached by SQL (and you have not specified the option `--nosql`), the script starts collecting diagnosis information. If the system cannot be reached by SQL, the script starts collecting support information but does not export data from system views.

Results

The script creates a zip file containing the collected information and saves it to the directory `DIR_GLOBAL/sapcontrol/snapshots`. `DIR_GLOBAL` typically points to `/usr/sap/<sid>/SYS/global`.

The name of the zip file is structured as follows:

```
fullsysteminfodump_<SID>_<DBNAME>_<HOST>_<timestamp>.zip
```

The timestamp in the file name is UTC. The host and SID are taken from the `sapprofile.ini` file.

The output directory for the zip file is shown as console output when the script is running.

Related Information

[Collect and Download Diagnosis Information with the Cockpit \[page 687\]](#)

[Collect and Download Diagnosis Information in SAP HANA Studio \[page 689\]](#)

[fullSystemInfoDump.py Command Line Options \[page 692\]](#)

[Diagnosis Information Collected \[page 694\]](#)

6.11.4.4 fullSystemInfoDump.py Command Line Options

You can specify several command line options when executing the `fullSystemInfoDump.py` script from the command line.

Option	Description
<code>--version</code>	Displays script version number
<code>--help</code>	Shows help
<code>--nosql</code>	Excludes the collection of system views

i Note

If you are connected to the system database, only information from the system views of the system database will be collected. Information from the system views of tenant databases will **not** be collected regardless of this option.

Option	Description
<code>--file <filename></code>	<p>Zips the specified file in its source directory</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 5px;"> <p>i Note</p> <p>This option only zips the file; it does not trigger the collection of any other information.</p> </div>
<code>--days <no. of days></code>	<p>Collects information from the specified number of past days</p> <p>The default value is 7.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 5px;"> <p>i Note</p> <p>You cannot use this option with the options <code>--fromDate</code> and <code>--toDate</code>.</p> </div>
<code>--fromDate <YYYY-MM-DD></code>	Collects information starting from the specified date
<code>--toDate <YYYY-MM-DD></code>	Collects information up to the specified date
<code>--rtedump</code>	<p>Restricts the information collected to an RTE dump file or files</p> <p>You can configure the creation and collection of RTE dump files further with the remaining options.</p>
<code>--indexservers <comma-separated list of index servers></code>	<p>Specifies the index server(s) from which RTE dump files are to be collected</p> <p>By default, dump files are created and collected for all index servers</p>
<code>--interval <interval in minutes></code>	<p>Specifies the interval at which RTE dump files are to be collected</p> <p>Possible values are 1, 5, 10, 15, and 30. The default value is 1.</p>
<code>--sets <no. of RTE dump file sets></code>	<p>Specifies the number of RTE dump file sets to be collected.</p> <p>Possible values are 1, 2, 3, 4, and 5.</p>
<code>--tenant <tenant database name></code>	<p>Specifies which tenant database information is to be collected from</p> <p>You must specify a database name.</p> <p>To collect information from the system database, specify SYSTEMDB.</p>

6.11.4.5 Diagnosis Information Collected

The Python support script `fullSystemInfoDump.py` script collects a range of information from your system for diagnosis purposes. It can be triggered from the SAP HANA cockpit, the SAP HANA studio, or directly from the command line.

i Note

All of the following file types are collected unless the option `--rtedump` is specified, in which case only runtime environment (RTE) dump files are created and collected.

Log File

All information about what has been collected is shown as console output and is written to a file named `log.txt` that is stored in the zip file.

Trace Files

Each of the following trace files is put into a file with the same name as the trace file. For storage reasons, only the trace files from the last 7 days are collected unabridged. Older trace files are not collected. This behavior can be changed by using option `--days` or with the options `--fromDate` and `--toDate`.

Crashdump files and runtime dump files are always collected unabridged.

- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/compileserver_alert_<SAPLOCALHOST>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/compileserver_<SAPLOCALHOST>.<...>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/daemon_<SAPLOCALHOST>.<...>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/indexserver_alert_<SAPLOCALHOST>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/indexserver_<SAPLOCALHOST>.<...>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/nameserver_alert_<SAPLOCALHOST>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/nameserver_history.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/nameserver_<SAPLOCALHOST>.<...>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/preprocessor_alert_<SAPLOCALHOST>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/preprocessor_<SAPLOCALHOST>.<...>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/statisticsserver_alert_<SAPLOCALHOST>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/statisticsserver_<SAPLOCALHOST>.<...>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/xsengine_alert_<SAPLOCALHOST>.trc`
- `$DIR_INSTANCE/<SAPLOCALHOST>/trace/xsengine_<SAPLOCALHOST>.<...>.trc`

Configuration Files

All configuration files are collected unabridged and stored in a file with the same name as the .ini file:

- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/attributes.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/compileserver.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/daemon.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/executor.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/extensions.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/filter.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/global.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/indexserver.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/inifiles.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/localclient.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/mimetypermapping.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/nameserver.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/preprocessor.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/scriptserver.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/statisticsserver.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/validmimetypes.ini
- \$DIR_INSTANCE/<SAPLOCALHOST>/exe/config/xsengine.ini

Database System Log Files

The following backup files are collected unabridged:

- \$DIR_INSTANCE/<SAPLOCALHOST>/trace/backup.log
- \$DIR_INSTANCE/<SAPLOCALHOST>/trace/backint.log

RTE Dump Files

For each index server, an RTE dump file containing information about threads, stack contexts, and so on is created and stored in the file `indexserver_<SAPLOCALHOST>_<PORT>_runtimedump.trc`. These files are stored unabridged.

Crashdump Information

Crashdump files for services are collected unabridged.

Performance Trace Files

Performance trace files with the suffix *.tpt are collected unabridged.

Kerberos Files

The following Kerberos files are collected:

- /etc/krb5.conf
- /etc/krb5.keytab

System Views

If the collection of system views is not excluded (option --nosql specified), all rows of the following system views (with the exceptions mentioned below) are exported into a CSV file with the name of the table.

i Note

If you are connected to the system database of a multiple-container system, only information from the system views of the system database will be collected. Information from the system views of tenant databases will **not** be collected regardless of this option.

- SYS.M_CE_CALCSCENARIOS WHERE SCENARIO_NAME LIKE '%_SYS_PLE%'
- SYS.M_CONNECTIONS with CONNECTION_ID > 0
- SYS.M_DATABASE_HISTORY
- SYS.M_DEV_ALL_LICENSES
- SYS.M_DEV_PLE_SESSIONS_
- SYS.M_DEV_PLE_RUNTIME_OBJECTS_
- SYS.M_EPM_SESSIONS
- SYS.M_INIFILE_CONTENTS
- SYS.M_LANDSCAPE_HOST_CONFIGURATION
- SYS.M_RECORD_LOCKS
- SYS.M_SERVICE_STATISTICS
- SYS.M_SERVICE_THREADS
- SYS.M_SYSTEM_OVERVIEW
- SYS.M_TABLE_LOCATIONS
- SYS.M_TABLE_LOCKS
- SYS.M_TABLE_TRANSACTIONS
- _SYS_EPM.VERSIONS
- _SYS_EPM.TEMPORARY_CONTAINERS
- _SYS_EPM.SAVED_CONTAINERS
- _SYS_STATISTICS.STATISTICS_ALERT_INFORMATION

- `_SYS_STATISTICS.STATISTICS_ALERT_LAST_CHECK_INFORMATION`

i Note

Only the first 2,000 rows are exported.

- `_SYS_STATISTICS.STATISTICS_ALERTS`

i Note

Only the first 2,000 rows are exported.

- `_SYS_STATISTICS.STATISTICS_INTERVAL_INFORMATION`
- `_SYS_STATISTICS.STATISTICS_LASTVALUES`
- `_SYS_STATISTICS.STATISTICS_STATE`
- `_SYS_STATISTICS.STATISTICS_VERSION`

The first 2,000 rows of all remaining tables in schema `_SYS_STATISTICS` are exported ordered by column `SNAPSHOT_ID`.

Additional Information Collected If SQL Connection Is Not Available

All available topology information is exported to a file named `topology.txt`. It contains information about the host topology in a tree-like structure. The keys are grouped using brackets while the corresponding values are referenced by the symbol `==>`. For example:

```
[
  ['host']
    ['host', 'ld8521']
      ['host', 'ld8521', 'role']
        ==> worker
      ['host', 'ld8521', 'group']
        ==> default
      ['host', 'ld8521', 'nameserver']
        ['host', 'ld8521', 'nameserver', '30501']
          ['host', 'ld8521', 'nameserver', '30501', 'activated_at']
            ==> 2011-08-09 16:44:02.684
          ['host', 'ld8521', 'nameserver', '30501', 'active']
            ==> no
          ['host', 'ld8521', 'nameserver', '30501', 'info']
            ['host', 'ld8521', 'nameserver', '30501', 'info', 'cpu_manufacturer']
              ==> GenuineIntel
            ['host', 'ld8521', 'nameserver', '30501', 'info',
'topology_mem_type']
              ==> shared
            ['host', 'ld8521', 'nameserver', '30501', 'info',
'sap_retrieval_path_devid']
              ==> 29
            ['host', 'ld8521', 'nameserver', '30501', 'info', 'build_time']
              ==> 2011-07-26 17:15:05
            ['host', 'ld8521', 'nameserver', '30501', 'info', 'net_realhostname']
              ==> -
            ['host', 'ld8521', 'nameserver', '30501', 'info', 'build_branch']
              ==> orange_COR
            ['host', 'ld8521', 'nameserver', '30501', 'info', 'mem_swap']
              ==> 34359730176
            ['host', 'ld8521', 'nameserver', '30501', 'info', 'mem_phys']
```

6.11.5 Problem Analysis Using hdbcons

`hdbcons` is a command line tool with which commands can be executed against running processes using a separate communication channel. It is intended for problem analysis by SAP HANA development support.

⚠ Caution

Technical expertise is required to use `hdbcons`. To avoid incorrect usage, use `hdbcons` only with the guidance of SAP HANA development support.

`hdbcons` commands can be executed directly in the Administration editor on the *Console* tab. However, it is not visible by default. You can enable the display of the *Console* tab in the preferences of the SAP HANA studio under ► *SAP HANA* ► *Global Settings* 🗑.

To see a list of available commands and display the help for a command, enter the command `help`.

Each command is subject to an individual authorization check. Operating system user (<sid>adm) access is not required.

6.11.6 Open a Support Connection

In some support situations, it may be necessary to allow an SAP support engineer to log into your system to analyze the situation.

Procedure

1. To enable a support user to log on to your system, complete the following tasks:
 - a. Install the SAProuter as described on SAP Support Portal.
 - b. Set up a support connection as described in SAP Note 1634848 (*SAP HANA database service connections*).
 - c. Configure a Telnet connection as described in SAP Note 37001 (*Telnet link to customer systems*).
 - d. Configure an SAP HANA database connection as described in SAP Note 1592925 (*SAP HANA studio service connection*).
 - e. Configure a TREX/BIA/HANA service connection as described in SAP Note 1058533 (*TREX/BIA/HANA service connection to customer systems*).
2. Create a database user and grant the MONITORING role.

The MONITORING role allows a database user to open the SAP HANA Administration Console perspective of the SAP HANA studio with read-only access to the system, system views, statistics views, trace files, and so on. However, this role does not provide any privileges for accessing application data. With the MONITORING role, it is also not possible to change the configuration of or start and stop a system. You can grant the MONITORING role to a support engineer if SAP support needs to connect to the system. Depending on the issue to be analyzed, further privileges may be needed to allow sufficient analysis (for example, to access application data or data models).

Related Information

[SAProuter](#)

[SAP Note 1634848](#)

[SAP Note 37001](#)

[SAP Note 1592925](#)

[SAP Note 1058533](#)

7 Security Administration and User Management

Security administration, including user management, represents a category of administration usually handled separately from general system administration tasks.

Related Information

[Monitoring Critical Security Settings in SAP HANA Cockpit \[page 700\]](#)

[Managing SAP HANA Users \[page 706\]](#)

[Auditing Activity in the SAP HANA Database \[page 826\]](#)

[Managing Data Encryption in SAP HANA \[page 847\]](#)

[Managing Client Certificates \[page 900\]](#)

[Data Anonymization \[page 916\]](#)

7.1 Monitoring Critical Security Settings in SAP HANA Cockpit

SAP HANA has many configuration settings that allow you to customize your system for your implementation scenario and system environment. Some of these settings are important for the security of your system. Misconfiguration could leave your system vulnerable. The SAP HANA cockpit allows you to monitor several critical security settings at a glance.

i Note

In addition to using SAP HANA cockpit to monitor critical security settings, please refer to *SAP HANA Security Checklists and Recommendations*. This document provides more detailed information as well as recommendations for many settings.

[View Status of Security Settings \[page 701\]](#)

You can view the status of critical security settings of the SAP HANA database in the SAP HANA cockpit on the [Overview](#) page.

[Security Tiles and Links \[page 701\]](#)

The [Security](#) section of the [Overview](#) page contains information related to critical security settings and links to further information and configuration options.

[Network Security Details \[page 704\]](#)

You can view important configuration settings related to secure internal SAP HANA communication and secure external SQL client communication in the SAP HANA cockpit.

7.1.1 View Status of Security Settings

You can view the status of critical security settings of the SAP HANA database in the SAP HANA cockpit on the [Overview](#) page.

Prerequisites

You have the authorization to see security-related information as described in the section *Security Tiles and Links*.

Procedure

1. On the [Overview](#) page, navigate to the [Security](#) area.
2. Review the security status displayed on the various tiles, drilling down for more detailed information and functions.

Related Information

[Security Tiles and Links \[page 701\]](#)

7.1.2 Security Tiles and Links

The [Security](#) section of the [Overview](#) page contains information related to critical security settings and links to further information and configuration options.

The screenshot shows a 'Security' dashboard with four main sections:

- Data Encryption:** Contains three rows, each with a toggle switch set to 'ON' and a timestamp for when the root key was changed (e.g., '16 Jun 2017, 08:17:10').
- Auditing:** Shows 'Auditing Status' as 'ON' with a green checkmark. Below, it lists 'Audit Trail Targets' (Database table), 'Enabled Audit Policies' (1), and 'Disabled Audit Policies' (3).
- Authentication:** Lists 'Password Policy' (Customized), 'Single Sign-on' (Kerberos, SAML), and 'SYSTEM User Password' (Changed on 9 May 2017).
- Security Related Links:** A list of links including 'Manage certificates', 'Manage certificates collections', 'Network security information', 'Security administration help', 'SAP HANA security website', and 'Security checklists'.

Security Tiles and Links

Tiles

Tile	Description	Required Authorization
Data Encryption	<p>Indicates the status of data volume encryption, log volume encryption, and backup encryption</p> <p>The on/off switches allow to you enable or disable each type of encryption.</p> <div data-bbox="603 622 986 853" style="border: 1px solid orange; padding: 5px; margin: 10px 0;"> <p>⚠ Caution</p> <p>Do not enable data volume encryption in an existing operational database without having first read the section <i>Enabling and Disabling Encryption of Data and Log Volumes</i>.</p> </div> <p>If you are connected to the system database, you will also see when the master keys of the secure stores in the file system (SSFS) were changed.</p> <p>This tile opens the Data Encryption Configuration page where you can see more information about the encryption status, enable or disable encryption, and change encryption keys.</p>	<p>System privilege CATALOG READ, and ENCRYPTION ROOT KEY ADMIN (to enable/disable encryption), and RE-SOURCE ADMIN (to view SSFS master key information)</p>
Auditing	<p>Indicates whether or not auditing is enabled in the database, the number of audit policies, and the configured audit trail target</p> <p>The Auditing switch allows you enable or disable auditing in the database.</p> <p>If a firefighter policy is active in the database (that is, a policy that audits all the actions of a particular user), this is also indicated.</p> <p>This tile opens the Auditing page where you can see more detailed information about audit policies, as well as create new ones. You can also make changes to global auditing settings.</p>	<p>To see information about auditing status audit policies, you need the system privileges AUDIT ADMIN or CATALOG READ.</p> <p>To see information about audit trail targets, you need the system privilege INI-FILE ADMIN.</p>

Tile	Description	Required Authorization
Authentication	<p>Indicates the status of the password policy (default or customized), the user authentication mechanisms configured for single sign-on in the database, as well as when the password of the SYS-TEM user was last changed</p> <p>This tile opens the Password Policy and Blacklist page where you can see and edit the password policy and blacklist.</p>	<p>System privilege CATALOG READ</p> <p>To be able to see the password blacklist on opening the Password Policy and Blacklist page, you need SELECT privilege on _SYS_SYS_PASS-WORD_BLACKLIST (_SYS_SECURITY).</p>

Security Related Links

Link	Description	Authorization
Manage certificates	Opens the Certificates page where you can import certificates into the certificate store	CATALOG READ, SSL ADMIN, USER ADMIN, TRUST ADMIN or CERTIFICATE ADMIN
Manage certificate collections	Opens the Certificate Collections page where you can create and configure certificate collections	CATALOG READ, SSL ADMIN, USER ADMIN, TRUST ADMIN or CERTIFICATE ADMIN
View network security information	Opens the Network Security Information page where you can see more detailed information about network configuration.	CATALOG READ
View anonymization report	Opens the Anonymization page where you can view all calculation views with anonymization node views configured	CATALOG READ
Manage SAML identity providers	Opens the SAML Identity Provider page where you can see existing identity providers in the database and add new ones	USER ADMIN
Security administration help	Opens the SAP HANA documentation that describes those security administration tasks that you can perform using the SAP HANA cockpit	No additional authorization required
SAP HANA security website	Opens the SAP HANA security website	No additional authorization required
Security checklists	Opens the document Security Checklists and Recommendations on SAP Help Portal	No additional authorization required

i Note

The links for managing users and roles are in [Administration](#) area of the [Overview](#) page.

Related Information

[Enabling Encryption of Data and Log Volumes \[page 864\]](#)

7.1.3 Network Security Details

You can view important configuration settings related to secure internal SAP HANA communication and secure external SQL client communication in the SAP HANA cockpit.

Note

For more information about how to configure secure communication, see the *SAP HANA Security Guide*.

General Settings

Field	Description
Cryptographic Provider	The cryptographic service provider being used by the SAP HANA server
Maximum TLS/SSL Protocol Version Accepted	The maximum TLS/SSL protocol version accepted
Minimum TLS/SSL Protocol Version Accepted	The minimum TLS/SSL protocol version accepted
Allowed TLS/SSL Cipher Suites	The encryption algorithms allowed for TLS/SSL connections This value depends on the cryptographic service provider used. The default values are <i>PFS:HIGH::EC_HIGH:+EC_OPT</i> (CommonCryptoLib) and <i>ALL:!ADH:!LOW:!EXP:!NULL:@STRENGTH</i> (OpenSSL).

Internal Communication

Field	Description
TLS/SSL Secured	Indicates whether or not internal communication channels are secured using TLS/SSL The following values are possible: <ul style="list-style-type: none">• <i>Disabled</i> (default)• <i>System PKI</i>• <i>Manual configuration</i> For more information about these values, see <i>Server-Side TLS/SSL Configuration Properties for Internal Communication</i> in the <i>SAP HANA Security Guide</i> .

Field	Description
<i>Listening On</i>	<p>Indicates the listening interface for internal SAP HANA connections</p> <p>The following values are possible:</p> <ul style="list-style-type: none"> • <i>Local network</i> SAP HANA services listen on the loopback interface only (IP address 127.0.0.1). Only connections from the local machine are possible. This value is only relevant for single-host systems and is the recommended configuration. • <i>Global network</i> In multiple-host systems without a separate internal network, SAP HANA services listen on all available network interfaces. Connections from remote machines are possible. <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p>⚠ Caution</p> <p>This setting exposes the internal SAP HANA service ports. To avoid a vector for security attacks, it is strongly recommended that you secure SAP HANA internal ports with an additional firewall.</p> </div> <ul style="list-style-type: none"> • <i>Internal network</i> In multiple-host systems with a separate internal network, SAP HANA services listen on a network interface within the allowed network mask. Only connections from machines (hosts) in the internal network are possible. <p>For more information, see <i>Configuring the Network for Multiple Hosts</i> and <i>Configuring SAP HANA Inter-Service Communication</i> in the <i>SAP HANA Administration Guide</i>.</p>
<i>Internal Host Name Resolution</i>	<p>The IP addresses of the network adapters used for SAP HANA internal communication</p> <p>This is relevant for multiple-host systems with a separate internal network (service communication: <i>Internal network</i>).</p>
<i>Key Store</i>	The key store file that contains the server's private key(s)
<i>Trust Store</i>	The trust store file that contains the server's public certificate(s)
<i>Validate Client Certificates</i>	Indicates whether or not the certificate of the communication partner is validated

Field	Description
Enforce TLS/SSL for SQL Connections	Indicates whether all clients communicating with the SAP HANA database via the SQL interface are required to use a secured connection The database refuses SQL connection attempts that don't use TLS/SSL.
Key Store	The key store file that contains the server's private key(s)
Trust Store	The trust store file that contains the server's public certificate(s)
Validate Client Certificates	Indicates whether or not the certificate of the communication partner is validated

Related Information

[Configuring SAP HANA Inter-Service Communication \[page 1440\]](#)

[Configuring the Network for Multiple Hosts \[page 1438\]](#)

7.2 Managing SAP HANA Users

Every user who wants to work with the SAP HANA database must have a database user. As a user administrator, you create and provision the required users, as well as perform other tasks related to user administration.

Managing users in SAP HANA includes the following tasks:

- Configuring SAP HANA for the required user authentication mechanisms
- Creating roles and users
- Granting users the roles and privileges they require for their duties
- Other user administration tasks such as resolving authorization or authentication issues, deactivating users, and so on.

Note

Users of SAP HANA SAP HANA Extended Services (SAP HANA XS) advanced applications are managed independently of the SAP HANA database. Dedicated administration tools available for managing application users and roles. For more information, see the section on maintaining the SAP HANA XS advanced model runtime in the *SAP HANA Administration Guide*.

[Database Users \[page 707\]](#)

Every user who wants to work with the SAP HANA database must have a database user.

[Operating System User sidadm \[page 714\]](#)

The `<sid>adm` user is not a database user but a user at the operating system level. Also referred to as the operating system administrator, this user has unlimited access to all local resources related to SAP systems.

[User Authentication and Single-Sign On \[page 715\]](#)

The identity of database users accessing SAP HANA is verified through a process called authentication. SAP HANA supports several authentication mechanisms, several of which can be used for the integration of SAP HANA into single sign-on environments (SSO). The mechanisms used to authenticate individual users is specified as part of the user definition.

[User Authorization \[page 738\]](#)

After successful logon, the user's authorization to perform the requested operations on the requested objects is verified.

[Provisioning Users \[page 775\]](#)

As a user administrator, you create and configure database users, as well as authorize them to work with the SAP HANA database.

Related Information

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1647\]](#)

7.2.1 Database Users

Every user who wants to work with the SAP HANA database must have a database user.

Database users are created with either the `CREATE USER` or `CREATE RESTRICTED USER` statement, or using the SAP HANA cockpit.

Standard Users

Standard users correspond to users created with the `CREATE USER` statement. By default they can create objects in their own schema and read data in system views. Read access to system views is granted by the `PUBLIC` role, which is granted to every standard user.

Restricted Users

Restricted users, which are created with the `CREATE RESTRICTED USER` statement, initially have no privileges. Restricted users are intended for provisioning users who access SAP HANA through client applications and who are not intended to have full SQL access via an SQL console. If the privileges required to use the application are encapsulated within an application-specific role, then it is necessary to grant the user only this role. In this way, it can be ensured that users have only those privileges that are essential to their work.

Compared to standard database users, restricted users are initially limited in the following ways:

- They cannot create objects in the database as they are not authorized to create objects in their own database schema.
- They cannot view any data in the database as they are not granted the standard `PUBLIC` role.
- They are only able to connect to the database using HTTP/HTTPS.

For restricted users to connect via ODBC or JDBC, access for client connections must be enabled by executing the SQL statement `ALTER USER <user_name> ENABLE CLIENT CONNECT` or enabling the corresponding option for the user in the SAP HANA cockpit.

For full access to ODBC or JDBC functionality, users also require the predefined role `RESTRICTED_USER_ODBC_ACCESS` or `RESTRICTED_USER_JDBC_ACCESS`.

i Note

Disabling ODBC/JDBC access for a user, either a restricted user or a standard user, does not affect the user's authorizations or prevent the user from executing SQL commands via channels other than JDBC/ODBC. If the user has been granted SQL privileges (for example, system privileges and object privileges), he or she is still authorized to perform the corresponding database operations using, for example, a HTTP/HTTPS client.

A user administrator can convert a restricted user into a standard user (or vice versa) as follows:

- Granting (or revoking) the PUBLIC role
You can do this by editing the user in the SAP HANA cockpit or with the SQL statement `ALTER USER <username> GRANT | REVOKE ROLE PUBLIC`.
- Granting (or revoking) authorization to create objects in the user's own schema
You can do this by editing the user in the SAP HANA cockpit or with the SQL statement `ALTER USER <username> GRANT | REVOKE CREATE ANY ON OWN SCHEMA`.
- Enabling (or disabling) full SQL
You can do this by editing the user in the SAP HANA cockpit or with the SQL statement `ALTER USER <user_name> ENABLE CLIENT CONNECT`.

i Note

A user is only identified as a restricted user in system view `USERS` if he doesn't have the PUBLIC role or authorization for his own schema.

Predefined Database Users

When an SAP HANA database is created, several database users are created by default. The most important of these is the `SYSTEM` database user, which should be deactivated in production systems.

Several technical database users (that is, database users that do not correspond to real people) are also created, for example, `SYS` and `_SYS_REPO`.

For more information about other predefined database users, see the *SAP HANA Security Guide*.

[The SYSTEM User \[page 709\]](#)

The `SYSTEM` database user is the initial user that is created during the creation of the SAP HANA database.

[Deactivate the SYSTEM User \[page 710\]](#)

As the most powerful database user, `SYSTEM` is not intended for use in production systems. Use it to create lesser privileged users for particular purposes and then deactivate it.

[Resetting the SYSTEM User Password \[page 711\]](#)

The system database and all tenant databases each have their own SYSTEM user. The system administrator can reset the password of any SYSTEM user if it has been irretrievably lost.

7.2.1.1 The SYSTEM User

The SYSTEM database user is the initial user that is created during the creation of the SAP HANA database.

SYSTEM is the database superuser. It has irrevocable system privileges, such as the ability to create other database users, access system tables, and so on.

In the system database, the SYSTEM user has additional privileges for managing tenant databases, for example, creating and dropping databases, changing configuration (*.ini) files of databases, and performing database-specific data backups.

It is highly recommended that you do not use SYSTEM for day-to-day activities in production environments. Instead, use it to create database users with the minimum privilege set required for their duties (for example, user administration, system administration). Then deactivate SYSTEM. You may temporarily reactivate the SYSTEM user for emergency or bootstrapping tasks.

i Note

The SYSTEM user is not required to update the SAP HANA database system; a lesser-privileged user can be created for this purpose. However, to upgrade SAP support package stacks, SAP enhancement packages and SAP systems using the Software Update Manager (SUM) and to install, migrate, and provision SAP systems using the Software Provisioning Manager (SWPM), the SYSTEM user **is required** and needs to be temporarily reactivated for the duration of the upgrade, installation, migration or provisioning.

If the password of SYSTEM user of the system database is lost, it can be reset using the operating system user (<sid>adm user). The system administrator can reset the SYSTEM user password of a tenant database from the system database.

Related Information

[Deactivate the SYSTEM User \[page 710\]](#)

[Resetting the SYSTEM User Password \[page 711\]](#)

7.2.1.2 Deactivate the SYSTEM User

As the most powerful database user, `SYSTEM` is not intended for use in production systems. Use it to create lesser privileged users for particular purposes and then deactivate it.

Prerequisites

You have the system privilege `USER ADMIN`.

Context

It is highly recommended that you do not use `SYSTEM` for day-to-day activities in production environments. Instead, use it to create database users with the minimum privilege set required for their duties (for example, user administration, system administration). Then deactivate `SYSTEM`. You may temporarily reactivate the `SYSTEM` user for emergency or bootstrapping tasks.

i Note

The `SYSTEM` user is not required to update the SAP HANA database system; a lesser-privileged user can be created for this purpose. However, to upgrade SAP support package stacks, SAP enhancement packages and SAP systems using the Software Update Manager (SUM) and to install, migrate, and provision SAP systems using the Software Provisioning Manager (SWPM), the `SYSTEM` user **is required** and needs to be temporarily reactivated for the duration of the upgrade, installation, migration or provisioning.

Procedure

Execute the following statement:

```
ALTER USER SYSTEM DEACTIVATE USER NOW
```

Results

The `SYSTEM` user is deactivated and can no longer **connect** to the SAP HANA database. However, it may appear as though the `SYSTEM` user is still active in the system (for example when a procedure that was created by `SYSTEM` with `DEFINER MODE` is called).

You can verify that this is the case in the `USERS` system view. For user `SYSTEM`, check the values in the columns `USER_DEACTIVATED`, `DEACTIVATION_TIME`, and `LAST_SUCCESSFUL_CONNECT`.

i Note

You can still use the `SYSTEM` user as an emergency user even if it has been deactivated. Any user with the system privilege `USER ADMIN` can reactivate `SYSTEM` with the statement `ALTER USER SYSTEM`

`ACTIVATE USER NOW`. To ensure that an administrator does not do this surreptitiously, it is recommended that you create an audit policy monitoring `ALTER USER` statements. Also change the password of the `SYSTEM` user after reactivating it.

7.2.1.3 Resetting the SYSTEM User Password

The system database and all tenant databases each have their own `SYSTEM` user. The system administrator can reset the password of any `SYSTEM` user if it has been irretrievably lost.

Related Information

[Reset the SYSTEM User Password of the System Database \[page 711\]](#)

[Reset the SYSTEM User Password of a Tenant Database \[page 713\]](#)

7.2.1.3.1 Reset the SYSTEM User Password of the System Database

If the password of the `SYSTEM` user of the system database is lost, you can reset it as the operating system administrator by starting the name server in emergency mode.

Prerequisites

- You cannot log on to the database as the `SYSTEM` because the password has been irretrievably lost.

Note

If you can log on as `SYSTEM` and you want to change the password, do not use the procedure described here. Use the SAP HANA studio or execute the `ALTER USER` SQL statement directly: `ALTER USER SYSTEM PASSWORD <new_password>`.

- You have the credentials of the operating system administrator (`<sid>adm`).

Procedure

- Log on to the server on which the name server of the system database is running as the operating system user (that is, `<sid>adm` user).

2. Open a command line interface.
3. Shut down the instance by executing the following command:


```
/usr/sap/<SID>/HDB<instance>/exe/sapcontrol -nr <instance> -function StopSystem
HDB
```
4. In a new session, start the name server of the system database by executing the following commands:
 - /usr/sap/<SID>/HDB<instance>/hdbenv.sh
 - /usr/sap/<SID>/HDB<instance>/exe/hdbnameserver -resetUserSystem

After some start-up notifications, the prompt resetting of user SYSTEM - new password appears, followed by additional notifications:

```

:usr/sap/ /> /usr/sap/ /exe/hdbnameserver -resetUserSystem
em
Starting interactive mode for resetting user SYSTEM...
unclean shutdown of service instance with pid 29905.
service startup...
accepting requests at
searching for master nameserver ...
assign as master nameserver. assign to volume 1 started
service startup...
Checking for recovery request ...
Loading topology ...
Opening persistence ...
run as transaction master
Loading topology ...
Loading licensing ...
setStarting(nameserver@:)
setActive(nameserver@)
service assigned as master
service start as systemsserver
setInactive(preprocessor@)
setInactive(webdispatcher@)
setInactive(compileserver@)
setInactive(indexserver@)
resetting of user SYSTEM - new password:

HDB      HDBSettings.sh  hdbenv.csh      work/
HDBAdmin.sh  backup/         hdbenv.sh       xterms
HDBSettings.csh  exe/          /

HDB      HDBSettings.sh  hdbenv.csh      work/
HDBAdmin.sh  backup/         hdbenv.sh       xterms
HDBSettings.csh  exe/          /

HDB      HDBSettings.sh  hdbenv.csh      work/
HDBAdmin.sh  backup/         hdbenv.sh       xterms
HDBSettings.csh  exe/          /

HDB      HDBSettings.sh  hdbenv.csh      work/
HDBAdmin.sh  backup/         hdbenv.sh       xterms
HDBSettings.csh  exe/          /
NewPassword1
new pw accepted.
(Re)Activating user SYSTEM...
done

```

Reset SYSTEM User Password (System Database)

5. After the appearance of the last notification, enter a new password for the SYSTEM user.
You must enter a password that complies with the password policy configured for the system.

The password for the SYSTEM user of the system database is reset and the name server stops.

6. In a new session, start the instance by executing the following command:

```
/usr/sap/<SID>/HDB<instance>/exe/sapcontrol -nr <instance> -function StartSystem  
HDB
```

Results

The password of the SYSTEM user of the system database is reset. You have to change the new password the next time you log on with this user.

If you previously deactivated the SYSTEM user, it is now also reactivated. This means you will need to deactivate it again.

7.2.1.3.2 Reset the SYSTEM User Password of a Tenant Database

If the password of the SYSTEM user in a tenant database is lost, you as the system administrator can reset it from the system database.

Prerequisites

- You cannot log on to the database as the SYSTEM because the password has been irretrievably lost.
- There is no user available with the system privilege USER ADMIN who can reset the SYSTEM user password.

Note

If you can log on as SYSTEM or another user with the system privilege USER ADMIN, do not use the procedure described here to change the password of the SYSTEM user. Instead, log on to the tenant database directly and either execute the `ALTER USER SYSTEM PASSWORD <new_password>` statement directly or change the password using the *User* editor in the SAP HANA cockpit

- You are connected to the system database and have the system privilege DATABASE ADMIN.

Procedure

→ Tip

You can also reset the SYSTEM user password of a tenant database using the SAP HANA cockpit.

1. Stop the tenant database, for example by executing the following statement:

```
ALTER SYSTEM STOP DATABASE <database_name>
```

2. Create a new password for the SYSTEM user by executing the following statement:

```
ALTER DATABASE <database_name> SYSTEM USER PASSWORD <new_password>
```

Note

The password must adhere to the password policy of the system database.

The password for the SYSTEM user is reset and the tenant database is started.

Results

- The password of the SYSTEM user of the tenant database is reset. You have to change the password the next time you log on with this user, this time in line with the password policy of the tenant database.
- If the SYSTEM user was previously deactivated, locked, or expired, it is now activated again. We recommend that you deactivate it.
- If auditing is enabled, the password change is automatically logged in both the system and tenant database audit trails.

Related Information

[Deactivate the SYSTEM User \[page 710\]](#)

[Auditing Activity in the SAP HANA Database \[page 826\]](#)

[Change a Database User \[page 802\]](#)

[Reset the SYSTEM Password of a Tenant using the Cockpit \[page 242\]](#)

7.2.2 Operating System User <sid>adm

The <sid>adm user is not a database user but a user at the operating system level. Also referred to as the operating system administrator, this user has unlimited access to all local resources related to SAP systems.

As part of the installation process, an external operating system user (<sid>adm, for example, spladm or xyzadm) is created.

This operating system user exists to provide an operating system context. From the operating system perspective, the operating system administrator is the user that owns all SAP HANA files and all related operating system processes. Certain administration operations require the operating system user's credentials, for example, starting or stopping the system.

If the system is configured for high isolation, additional OS users must be created for each tenant database. As a result, the processes of individual tenant databases run under dedicated OS users belonging to dedicated OS groups, and not under <sid>adm. Database-specific data on the file system is subsequently protected using standard OS file and directory permissions.

Related Information

[Database Isolation \[page 204\]](#)

7.2.3 User Authentication and Single-Sign On

The identity of database users accessing SAP HANA is verified through a process called authentication. SAP HANA supports several authentication mechanisms, several of which can be used for the integration of SAP HANA into single sign-on environments (SSO). The mechanisms used to authenticate individual users is specified as part of the user definition.

i Note

For JDBC and ODBC client connections, user passwords are always transmitted in encrypted hashed form during the user authentication process, never in plain text. For HTTP connections via SAP HANA XS classic, HTTPS must be configured. In SSO environments, we recommend using encrypted communication channels for **all** client connections.

[User Authentication Mechanisms \[page 715\]](#)

Authentication mechanisms supported in SAP HANA. Mechanisms that are not required can be disabled.

[Configuring SAP HANA for User Authentication and Single-Sign On \[page 719\]](#)

You can integrate SAP HANA into the user authentication infrastructure of your system landscape. To do so, you must configure SAP HANA for the required mechanisms.

7.2.3.1 User Authentication Mechanisms

Authentication mechanisms supported in SAP HANA. Mechanisms that are not required can be disabled.

Supported Authentication Mechanisms

Mechanism	Description	Can Be Used for SSO
SAP HANA user name and password	Users accessing the SAP HANA database authenticate themselves by entering their database user name and their local SAP HANA password.	No

Mechanism	Description	Can Be Used for SSO
Kerberos, SPNEGO	<p>A Kerberos authentication provider can be used to authenticate users accessing SAP HANA in the following ways:</p> <ul style="list-style-type: none"> • Directly from ODBC and JDBC database clients within a network (for example, the SAP HANA studio) • Indirectly from front-end applications such as SAP BusinessObjects applications and other SAP HANA databases using Kerberos delegation • Via HTTP/HTTPS access by means of SAP HANA Extended Services (SAP HANA XS), advanced model and classic model <p>In this case, Kerberos authentication is enabled with Simple and Protected GSSAPI Negotiation Mechanism (SPNEGO).</p>	Yes
<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>A user who connects to the database using an external authentication provider must also have a database user known to the database. The external identity is mapped to the identity of an internal database user.</p> </div>		
Security assertion markup language (SAML)	<p>A SAML bearer assertion can be used to authenticate users accessing SAP HANA directly from ODBC/JDBC database clients. SAP HANA can act as a service provider to authenticate users accessing via HTTP/HTTPS by means of SAP HANA XS classic and advanced.</p>	Yes
<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>A user who connects to the database using an external authentication provider must also have a database user known to the database. The external identity is mapped to the identity of an internal database user.</p> </div>		
Logon and assertion tickets	<p>Users can be authenticated by SAP logon or assertion tickets issued to them when they log on to an SAP system that is configured to create tickets (for example, the SAP Web Application Server or Portal).</p>	Yes
<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>To implement logon/assertion tickets, the user specified in the logon/assertion ticket must already exist in SAP HANA; there is no support for user mapping.</p> </div>		

Mechanism	Description	Can Be Used for SSO
X.509 client certificates	<p>For HTTP/HTTPS access to SAP HANA by means of SAP HANA XS advanced model and classic model, users can be authenticated by client certificates signed by a trusted Certification Authority (CA), which can be stored in the SAP HANA XS trust store.</p> <p>i Note To implement X.509 client certificates, the user specified in the certificate must already exist in SAP HANA; there is no support for user mapping.</p>	Yes for HTTP/HTTPS access to SAP HANA by means of SAP HANA XS (advanced and classic)
JSON Web Token (JWT)	<p>A JSON Web Token can be used to authenticate users accessing SAP HANA directly from ODBC/JDBC database clients or indirectly through SAP HANA extended application services, advanced model (SAP HANA XS, advanced).</p> <p>i Note A user who connects to the database using an external authentication provider must also have a database user known to the database. The external identity is mapped to the identity of an internal database user.</p>	Yes
LDAP	<p>A password stored in an LDAP directory server can be used to authenticate users accessing SAP HANA directly from ODBC/JDBC database clients, if authentication using users' local SAP HANA authentication has been disabled.</p> <p>i Note A user who connects to the database using an external authentication provider must also have a database user known to the database. The external identity and the database user name are the same. If the LDAP provider is enabled to create database users in SAP HANA, the required user is created automatically if it doesn't exist.</p>	No
Session cookies	<p>Session cookies are not technically an authentication mechanism. However, they reconnect users who have already been authenticated by Kerberos or SAML and extend the validity period of logon and assertion tickets.</p>	Yes

Isolated Single Sign-On for Tenant Databases

Separate, database-specific authentication is possible for every certificate-based authentication mechanism since it is possible to create different certificate collections for individual purposes directly in every database, and every database can have its own key pair and public key certificate.

For SAML assertions, X.509 certificates, JSON web tokens, and logon tickets, it is also possible to use certificate collections (or PSEs) located on the file system. It is still possible to configure different trust and key stores for every database in the `global.ini` file. However, bear the following points in mind:

- If different trust and key stores are not explicitly configured for tenant databases, the same ones will be used for all external communication channels (including HTTP) for all databases.

⚠ Caution

If you have configured in tenant databases or the system database single sign-on mechanisms that rely on trust stores located in the file system (such as SAP logon and assertion tickets or SAML) and the trust stores are shared, users of one tenant database may be able to log on to other databases in the system.

- By default, only the system administrator can configure separate trust and key stores for tenant databases by changing the relevant properties in the `global.ini` file. This is because tenant database administrators are prevented from changing any communication properties. They are in the default configuration change blacklist (`multidb.ini`).

For more information about certificate collections in the database and PSEs in the file system, see the section on certificate management.

For Kerberos-based authentication, a per-database configuration is not possible – databases users in all databases must be mapped to users in the same Key Distribution Center.

Disabling Authentication Mechanisms

By default all authentication mechanisms are enabled, but it is possible and recommended to disable those that are not used in your environment. You do this by configuring the parameter `[authentication] authentication_methods` in the `global.ini` configuration file. The value of this parameter specifies all enabled methods as a comma-separated list.

The default value is `password,kerberos,spnego,saml,saplogon,x509xs,jwt,sessioncookie,ldap`.

i Note

If you are using SAP HANA dynamic tiering, it is not possible to disable logon and assertion tickets (`saplogon`) as an authentication mechanism.

Changes to this parameter are audited by default if auditing is enabled.

Related Information

[Maintaining Single Sign-On for XS Advanced Applications \[page 1845\]](#)

7.2.3.2 Configuring SAP HANA for User Authentication and Single-Sign On

You can integrate SAP HANA into the user authentication infrastructure of your system landscape. To do so, you must configure SAP HANA for the required mechanisms.

SAP HANA supports several authentication mechanisms, several of which can be used for the integration of SAP HANA into single sign-on environments (SSO). Depending on which mechanisms you are implementing, you must configure SAP HANA accordingly.

Configuration of Authentication of SAP HANA XS Applications

To configure security-related aspects of SAP HANA XS applications, including SSO, use the SAP HANA XS Classic Administration Tools. For more information, see the section on maintaining the SAP HANA XS classic model run-time in the *SAP HANA Administration Guide*.

[Configure the Database Password Policy and Password Blacklist \[page 720\]](#)

The passwords of database users are subject to certain rules. These are defined in the password policy and the password blacklist. You can change the default password policy of the database and maintain entries in the password blacklist in line with your organization's security requirements.

[Configure Kerberos for SAP HANA Database Hosts \[page 730\]](#)

If you are implementing Kerberos-based user authentication, you must configure Kerberos on the authentication server.

[Add a SAML Identity Provider in SAP HANA Studio \[page 732\]](#)

If you are implementing Security Assertion Markup Language (SAML) to authenticate users accessing SAP HANA via the SQL interface directly (that is using JDBC and ODBC clients), you must add the SAML identity providers for the required users. You can do this using the SAP HANA studio.

[Add a SAML Identity Provider in SAP HANA Cockpit \[page 733\]](#)

If you are implementing Security Assertion Markup Language (SAML) to authenticate users accessing SAP HANA via the SQL interface directly (that is using JDBC and ODBC clients), you must add the SAML identity providers for the required users. You can do this using the SAP HANA cockpit.

[Configure the Trust Store for SAP Logon Tickets and Assertions \[page 735\]](#)

If you are integrating SAP HANA system into a landscape that uses SAP logon or assertion tickets for user authentication, you must configure SAP HANA to accept logon/assertion tickets.

[Configure an LDAP Server Connection for LDAP User Authentication \[page 736\]](#)

To enable LDAP user authentication, you set up a connection to an LDAP server by creating an LDAP provider in the SAP HANA database. Depending on your requirements, you configure the LDAP server to authenticate users only, or to authenticate and authorize users. For LDAP-authenticated users, you can also enable the automatic creation of users in SAP HANA.

Related Information

[Maintaining Single Sign-On for SAP HANA XS Applications \(Classic\) \[page 1583\]](#)

7.2.3.2.1 Configure the Database Password Policy and Password Blacklist

The passwords of database users are subject to certain rules. These are defined in the password policy and the password blacklist. You can change the default password policy of the database and maintain entries in the password blacklist in line with your organization's security requirements.

Context

The password policy of the database is defined by parameters in the `password_policy` section of the `indexserver.ini` configuration file for tenant databases and the `nameserver.ini` configuration file for the system database. The database password policy is valid for all database users unless the user is in a user group with its own dedicated password policy. For more information about user group-specific password policies, see the section on user groups in the *SAP HANA Security Guide*.

→ Tip

To determine which password policy a user is currently subject to, query the system view `M_EFFECTIVE_PASSWORD_POLICY`.

In addition to configuring the password policy parameters, you can also add words or partial words to the password blacklist. The password blacklist is implemented with the database table `_SYS_PASSWORD_BLACKLIST` in the schema `_SYS_SECURITY`. This table is empty when on database creation.

You can change the database password policy and edit the password blacklist using the *Password Policy and Blacklist* app of the SAP HANA cockpit or the *Security* editor of the SAP HANA studio.

Related Information

[Configure the Database Password Policy and Blacklist in SAP HANA Studio \[page 722\]](#)

[Configure the Database Password Policy and Blacklist in SAP HANA Cockpit \[page 721\]](#)

[Password Policy Configuration Options \[page 723\]](#)

7.2.3.2.1.1 Configure the Database Password Policy and Blacklist in SAP HANA Cockpit

Configure the password policy and password blacklist using the SAP HANA cockpit.

Prerequisites

- You have the system privilege INIFILE ADMIN.
- You have the object privileges SELECT, INSERT, and DELETE for the `_SYS_PASSWORD_BLACKLIST` table in the `_SYS_SECURITY` schema.

Procedure

1. In the SAP HANA cockpit, navigate to the *Overview* page and choose the *Authentication* block.

The *Password Policy and Blacklist* page opens.

2. Click *Edit* in the footer bar.

3. In the *Password Policy* area, configure the options in line with your security requirements.

All options have a default value. For more information about the individual parameters and their default values, see *Password Policy Configuration Options*.

4. In the *Password Blacklist* area, add the words or partial words that you want to prohibit in passwords.

The following configuration options are available:

Option	Description
Contained in Password	If you select this option, passwords that contain the blacklisted word are excluded. If you do not select this option, only passwords that match the blacklisted word exactly are excluded.
Case-Sensitive	If you select this option, the blacklisted word is case sensitive.

Example

If you add the words `SAP`, `my_sap_pwd`, and `sap_password` to the blacklist and select the *Contained in Password* checkbox, then passwords containing "SAP", "my_sap_pwd", and "sap_password" are not allowed, regardless of how the password policy is configured.

5. Click *Save* to save the password policy and password blacklist.

Results

The passwords of database users must be created and changed in line with the defined policy.

Related Information

[Password Policy Configuration Options \[page 723\]](#)

7.2.3.2.1.2 Configure the Database Password Policy and Blacklist in SAP HANA Studio

Configure the password policy and password blacklist using the SAP HANA studio.

Prerequisites

- You have the system privilege INIFILE ADMIN.
- You have the object privileges SELECT, INSERT, and DELETE for the `_SYS_PASSWORD_BLACKLIST` table (`_SYS_SECURITY`).

Procedure

1. Open the SAP HANA studio.
2. Open the *Security* editor of the database whose password policy you want to configure and choose the *Password Policy* tab.
3. In the *Password Policy* area, configure the options in line with your security requirements.
All options have a default value. For more information about the individual parameters and their default values, see *Password Policy Configuration Options*.
4. In the *Password Blacklist* area, add words or partial words that you want to prohibit in passwords by choosing the  (*Add*) button and entering the word.

The following configuration options are available:

Option	Description
Contained in Password	If you select this option, passwords that contain the blacklisted word are excluded. If you do not select this option, only passwords that match the blacklisted word exactly are excluded.
Case-Sensitive	If you select this option, the blacklisted word is case sensitive.

Example

If you add the words `SAP`, `my_sap_pwd`, and `sap_password` to the blacklist and select the *Contained in Password* checkbox, then passwords containing "SAP", "my_sap_pwd", and "sap_password" are not allowed, regardless of how the password policy is configured.

5. Choose the  (*Deploy*) button.

Results

The passwords of database users must be created and changed in line with the defined policy.

Related Information

[Password Policy Configuration Options \[page 723\]](#)

7.2.3.2.1.3 Password Policy Configuration Options

The password policy is defined by parameters in the `password_policy` section of the `indexserver.ini` configuration file (tenant databases) or `nameserver.ini` configuration file (system database). Password policy parameters may also be individually configured in the definition of a user group.

The following sections describe these parameters, which correspond to the configuration options available in the SAP HANA cockpit and SAP HANA studio.

- [Minimum Password Length \[page 723\]](#)
- [Lowercase Letters/Uppercase Letters/Numerical Digits/Special Characters Required \[page 724\]](#)
- [Password Change Required on First Logon \[page 725\]](#)
- [Number of Last Used Passwords That Cannot Be Reused \[page 725\]](#)
- [Number of Allowed Failed Logon Attempts \[page 726\]](#)
- [User Lock Time \[page 726\]](#)
- [Minimum Password Lifetime \[page 727\]](#)
- [Maximum Password Lifetime \[page 727\]](#)
- [Lifetime of Initial Password \[page 728\]](#)
- [Maximum Duration of User Inactivity \[page 728\]](#)
- [Notification of Password Expiration \[page 728\]](#)
- [Exempt SYSTEM User from Locking \[page 729\]](#)
- [Detailed Error Information on Failed Logon \[page 729\]](#)

Minimum Password Length

The minimum number of characters that the password must contain

Parameter	<code>minimal_password_length</code>
Default Value	8 (characters)
Additional Information	You must enter a value between 6 and 64.
UI Label	<i>Minimum Password Length</i>

Lowercase Letters/Uppercase Letters/Numerical Digits/ Special Characters Required

The character types that the password must contain and how many

Parameter	password_layout
Default Value	Aa1, that is, at least one uppercase letter, at least one number, and at least one lowercase letter
Additional Information	<p>The following character types are possible:</p> <ul style="list-style-type: none">• Lowercase letter (a-z)• Uppercase letter (A-Z)• Numerical digits (0-9)• Special characters (underscore (_), hyphen (-), and so on) Any character that is not an uppercase letter, a lowercase letter, or a numerical digit is considered a special character. <p>The following formats are supported for passwords:</p> <pre><password> ::= { <letter> [{ <letter_or_digit> # \$ } [...]] <digit> [<letter_or_digit> [...]] <any_quoted_string> }</pre> <p>If configuring this option in the <code>indexserver.ini</code> file using the <code>password_layout</code> parameter, you can use any specific letters, numbers and special characters, and the characters can be in any order. For example, the default value example could also be represented by a1A, hQ5, or 9fG. To enforce the use of at least one of each character type including special characters, you specify A1a_ or 2Bg?. To enforce the use of a specific number of a particular character type, specify the character type multiple times. For example, if passwords must contain at least 3 digits, you could specify the layout with a123A or 789fG.</p> <div data-bbox="603 1429 1396 1630"><p>Note</p><p>Passwords containing special characters other than underscore must be enclosed in double quotes ("). The SAP HANA studio does this automatically. When a password is enclosed in double quotes ("), any Unicode characters may be used.</p></div> <div data-bbox="603 1646 1396 1818"><p>Caution</p><p>The use of passwords enclosed in double quotes (") may cause logon issues depending on the client used. The SAP HANA studio and <code>hdsql</code> support passwords enclosed in double quotes (").</p></div>
UI Labels	<i>Lowercase Letters/Uppercase Letters/Numerical Digits/Special Characters Required</i>

Password Change Required on First Logon

Defines whether users have to change their initial passwords immediately the first time they log on

Parameter	<code>force_first_password_change</code>
Default Value	True
Additional Information	<p>If this parameter is set to true, users can still log on with the initial password but every action they try to perform will return the error message that they must change their password.</p> <p>If this parameter is set to false, users are not forced to change their initial password immediately the first time they log on. However, if a user does not change the password before the number of days specified in the parameter <code>maximum_unused_initial_password_lifetime</code>, then the password still expires and must be reset by a user administrator.</p> <p>A user administrator (that is, a user with the system privilege USER ADMIN) can force a user to change his or her password at any time with the following SQL statement: <code>ALTER USER <user_name> FORCE PASSWORD CHANGE</code></p> <p>A user administrator can override this password policy setting for individual users (for example, technical users) with the following SQL statement:</p> <ul style="list-style-type: none">• <code>CREATE USER <user_name> PASSWORD <password> [NO FORCE_FIRST_PASSWORD_CHANGE]</code>• <code>ALTER USER <user_name> PASSWORD <password> [NO FORCE_FIRST_PASSWORD_CHANGE]</code>
UI Label	<i>Password Change Required on First Logon</i>

Note

This parameter is only valid for users connecting with their SAP HANA database user name and password. It is not valid for connections established through other authentication mechanisms.

Number of Last Used Passwords That Cannot Be Reused

The number of last used passwords that the user is not allowed to reuse when changing his or her current password

Parameter	<code>last_used_passwords</code>
Default Value	5 (previous passwords)
Additional Information	If you enter the value 0 , the user can reuse his or her old password.
UI Label	<i>Number of Last Used Passwords That Cannot Be Reused</i>

Number of Allowed Failed Logon Attempts

The maximum number of failed logon attempts that are possible; the user is locked as soon as this number is reached

Parameter	<code>maximum_invalid_connect_attempts</code>
Default Value	6 (failed logon attempts)
Additional Information	<p>You must enter a value of at least 1.</p> <p>A user administrator can reset the number of invalid logon attempts with the following SQL statement: <code>ALTER USER <user_name> RESET CONNECT ATTEMPTS</code></p> <p>The first time a user logs on successfully after an invalid logon attempt, an entry is made in the <code>INVALID_CONNECT_ATTEMPTS</code> system view containing the following information:</p> <ul style="list-style-type: none">• The number of invalid logon attempts since the last successful logon• The time of the last successful logon <p>A user administrator can delete information about invalid logon attempts with the following SQL statement: <code>ALTER USER <user_name> DROP CONNECT ATTEMPTS</code></p> <div style="background-color: #f0f0f0; padding: 10px;"><p>→ Recommendation</p><p>Create an audit policy to log activity in the <code>INVALID_CONNECT_ATTEMPTS</code> system view. For example, create an audit policy that logs data query and manipulation statements executed on this view.</p></div> <div style="background-color: #f0f0f0; padding: 10px;"><p>i Note</p><p>Although this parameter is not valid for the <code>SYSTEM</code> user, the <code>SYSTEM</code> user will still be locked if the parameter <code>password_lock_for_system_user</code> is set to true. If <code>password_lock_for_system_user</code> is set to false, the <code>SYSTEM</code> user will not be locked regardless of the number of failed logon attempts.</p></div>
UI Label	<i>Number of Allowed Failed Logon Attempts</i>

User Lock Time

The number of minutes for which a user is locked after the maximum number of failed logon attempts

Parameter	<code>password_lock_time</code>
Default Value	1440 (minutes)

Additional Information

If you enter the value **0**, the user is unlocked immediately. This disables the functionality of parameter `maximum_invalid_connect_attempts`.

A user administrator can reset the number of invalid logon attempts and reactivate the user account with the following SQL statement: `ALTER USER <user_name> RESET CONNECT ATTEMPTS`. It is also possible to reactivate the user in the user editor of the SAP HANA Studio.

To lock a user indefinitely, enter the value **-1**. On the [Password Policy and Blacklist](#) page of the SAP HANA cockpit or in the [Security](#) editor of the SAP HANA studio, this corresponds to selecting the [Lock User Indefinitely](#) checkbox. The user remains locked until reactivated by a user administrator as described above.

UI Label	User Lock Time
----------	--------------------------------

Minimum Password Lifetime

The minimum number of days that must elapse before a user can change his or her password

Parameter	<code>minimum_password_lifetime</code>
Default Value	1 (day)
Additional Information	If you enter the value 0 , the password has no minimum lifetime.
UI Label	Minimum Password Lifetime

Maximum Password Lifetime

The number of days after which a user's password expires

Parameter	<code>maximum_password_lifetime</code>
Default Value	182 (days)
Additional Information	<p>You must enter a value of at least 1.</p> <p>A user administrator can exclude users from this password check with the following SQL statement: <code>ALTER USER <user_name> DISABLE PASSWORD LIFETIME</code>. However, this is recommended only for technical users only, not database users that correspond to real people.</p> <p>A user administrator can re-enable the password lifetime check for a user with the following SQL statement: <code>ALTER USER <user_name> ENABLE PASSWORD LIFETIME</code>.</p>
UI Label	Maximum Password Lifetime

Lifetime of Initial Password

The number of days for which the initial password or any password set by a user administrator for a user is valid

Parameter	<code>maximum_unused_initial_password_lifetime</code>
Default Value	7 (days)
Additional Information	<p>You must enter a value of at least 1.</p> <p>If a user has not logged on using the initial password within the given period of time, the user will be deactivated until their password is reset.</p>
	<div style="border-left: 2px solid #0070C0; padding-left: 10px;"><p>i Note</p><p>In SAP HANA 1.0 SPS 12 and earlier, this parameter was misspelled as <code>maximum_unused_inital_password_lifetime</code>. If this parameter had a user-specified value before upgrade, this value will be set as the value of the parameter <code>maximum_unused_initial_password_lifetime</code>. The misspelled parameter is unset and disappears from the custom configuration file.</p></div>
UI Label	<i>Lifetime of Initial Password</i>

Maximum Duration of User Inactivity

The number of days after which a password expires if the user has not logged on

Parameter	<code>maximum_unused_productive_password_lifetime</code>
Default Value	365 (days)
Additional Information	<p>You must enter a value of at least 1.</p> <p>If a user has not logged on within the given period of time using any authentication method, the user will be deactivated until their password is reset.</p>
UI Label	<i>Maximum Duration of User Inactivity</i>

Notification of Password Expiration

The number of days before a password is due to expire that the user receives notification

Parameter	<code>password_expire_warning_time</code>
Default Value	14 (days)

Additional Information

Notification is transmitted via the database client (ODBC or JDBC) and it is up to the client application to provide this information to the user.

If you enter the value **0**, the user does not receive notification that his or her password is due to expire.

The system also monitors when user passwords are due to expire and issues a medium priority alert (check 62). This may be useful for technical database users since password expiration results in the user being locked, which may affect application availability. It is recommended that you disable the password lifetime check of technical users so that their password never expires (`ALTER USER <technical_username> DISABLE PASSWORD LIFETIME`).

UI Label	<i>Notification of Password Expiration</i>
-----------------	--

Exempt SYSTEM User from Locking

Indicates whether or not the user SYSTEM is locked for the specified lock time (`password_lock_time`) after the maximum number of failed logon attempts (`maximum_invalid_connect_attempts`)

Parameter	<code>password_lock_for_system_user</code>
Default Value	true
Additional Information	This parameter cannot be configured for a user group.
UI Label	<i>Exempt SYSTEM User from Locking</i>

Detailed Error Information on Failed Logon

Indicates the detail level of error information returned when a logon attempt fails

Parameter	<code>detailed_error_on_connect</code>
Default Value	false
Additional Information	<p>If set to false, only the information <code>authentication failed</code> is returned.</p> <p>If set to true, the specific reason for failed logon is returned:</p> <ul style="list-style-type: none"> • Invalid user or password • User is locked • Connect try is outside validity period • User is deactivated
UI Label	<i>Detailed Error Information on Failed Logon</i>

Related Information

[Configure the Database Password Policy and Blacklist in SAP HANA Cockpit \[page 721\]](#)

[Configure the Database Password Policy and Blacklist in SAP HANA Studio \[page 722\]](#)

[Create an Audit Policy \[page 839\]](#)

7.2.3.2.2 Configure Kerberos for SAP HANA Database Hosts

If you are implementing Kerberos-based user authentication, you must configure Kerberos on the authentication server.

Prerequisites

To allow users to log on to the SAP HANA database using Kerberos authentication, you have installed MIT Kerberos client libraries on the host(s) of the SAP HANA database.

Context

SAP HANA supports Kerberos version 5 for single sign-on based on Active Directory (Microsoft Windows Server) or Kerberos authentication servers. For HTTP access via SAP HANA Extended Services (SAP HANA XS), advanced or classic model, Kerberos authentication is enabled with Simple and Protected GSSAPI Negotiation Mechanism (SPNEGO).

Once Kerberos client libraries have been installed, you must configure Kerberos on the authentication server by performing the following logical steps:

1. Register service principal names (SPN) for each host in the SAP HANA system using the following syntax:
`<service>/<host domain name>@<Kerberos realm name>`, where
 - `<service>` is either **hdb** (for Kerberos via ODBC and JDBC) or **HTTP** (for SPNEGO via HTTP/SAP HANA XS)
 - `<host domain name>` is the fully qualified domain name of the host
If the service is HTTP, you must register one SPN for each alias of the host name as well.
 - `<Kerberos realm name>` (Kerberos terminology) is identical to domain name in Active Directory terminology

This results in the generation of a service key table (keytab) for each host. This keytab contains the encrypted key for the host in question.

2. Export the keytab(s) to files.

3. Import each keytab file into the Kerberos installation on the respective host.

Procedure

The concrete steps to be performed on the authentication server depend on whether you are using Kerberos or Active Directory as follows:

1. Register the SPNs.

i Note

In Active Directory, before a SPN can be registered, you must create a plain user account that acts as the server principal on the domain controller. Afterward, you must map the SPN to the user account using a separate command.

2. Export the keytab(s) to files using a command line tool shipped with the authentication server.
This is applicable for both Kerberos and Active Directory.
3. Import the keytab files.
The files are transported to the file system path on the SAP HANA database hosts in line with how the Kerberos client is configured.

Results

You can now map the users stored in the Kerberos Key Distribution Center (KDC) to database users in SAP HANA database. You can do this when you create database users. Alternatively, if database users already exist, you can change their authentication details.

→ Remember

A per-database configuration is not possible – databases users in all databases must be mapped to users in the same KDC.

For more information about how to set up SSO with SAP HANA using Kerberos and Microsoft Active Directory, see SAP Note 1837331.

Related Information

[Maintaining the SAP HANA XS Classic Model Run Time \[page 1526\]](#)

[Create a Database User \[page 792\]](#)

[SAP Note 1837331](#)

7.2.3.2.3 Add a SAML Identity Provider in SAP HANA Studio

If you are implementing Security Assertion Markup Language (SAML) to authenticate users accessing SAP HANA via the SQL interface directly (that is using JDBC and ODBC clients), you must add the SAML identity providers for the required users. You can do this using the SAP HANA studio.

Prerequisites

- You have created a certificate collection with the purpose **SAML** in the database and have imported the X.509 certificates that will be used to sign the SAML assertions from the identity provider. Ensure that the entire certificate chain of the X.509 certificate is available.

⚠ Caution

We recommend creating certificate collections for individual purposes in the database directly, rather than using trust stores (PSE) in the file system. By default, the same PSE in the file system is shared by all databases for all external communication channels (including HTTP) and certificate-based authentication. Different PSEs must be explicitly configured for tenant databases.

- You have the system privilege USER ADMIN.

Procedure

i Note

While you can configure SAML providers for ODBC/JDBC-based SAML authentication using the SAP HANA cockpit, SAP HANA studio or SQL, always use the SAP HANA XS Administration Tool to configure SAML providers that will be used for HTTP access via the XS classic server.

- In the *Security* editor, choose the *SAML Identity Providers* tab.
- Select the relevant cryptographic provider.
- Add a new identity provider as follows:

Option	Description
Read from certificate	<ol style="list-style-type: none">Choose  (<i>Import SAML identity provider from certificate file</i>).Enter the name of the identity provider. The following naming conventions apply: Spaces and special characters except underscore () are not permitted. The name must start with a letter. The name cannot exceed 127 characters. <p>The system reads the X.509 certificate obtained from the identity provider and extracts the issuer and subject distinguished names (DNs). It then enters these in the corresponding fields.</p>

Option	Description
	<p>i Note</p> <p>If the certificate fails to read with an IOException or a CertificateException, try recoding the certificate from Base64 (*.pem) to DER (*.der) using OpenSSL or other tools.</p> <p>You can also enter the issuer and subject DNs manually.</p>
Manually	<ol style="list-style-type: none"> 1. Choose  (<i>Add SAML identity provider</i>) 2. Enter the name of the identity provider (in line with the above naming conventions). 3. Enter the issuer and subject DNs.

4. Save the identity provider by choosing the  (*Deploy*) button.

Results

The identity provider is now available for mapping to individual database users. You can do this when you create the database user. Alternatively, if the database user already exists, you can change their authentication details.

Related Information

[Managing Client Certificates \[page 900\]](#)

[Maintaining SAML Providers \(HTTP Access via XS Classic Server\) \[page 1559\]](#)

[Create a Database User \[page 792\]](#)

7.2.3.2.4 Add a SAML Identity Provider in SAP HANA Cockpit

If you are implementing Security Assertion Markup Language (SAML) to authenticate users accessing SAP HANA via the SQL interface directly (that is using JDBC and ODBC clients), you must add the SAML identity providers for the required users. You can do this using the SAP HANA cockpit.

Prerequisites

- You have created a certificate collection with the purpose **SAML** in the database and have imported the X.509 certificates that will be used to sign the SAML assertions from the identity provider. Ensure that the entire certificate chain of the X.509 certificate is available.

⚠ Caution

We recommend creating certificate collections for individual purposes in the database directly, rather than using trust stores (PSE) in the file system. By default, the same PSE in the file system is shared by all databases for all external communication channels (including HTTP) and certificate-based authentication. Different PSEs must be explicitly configured for tenant databases.

- You have the system privilege USER ADMIN.

Procedure

i Note

While you can configure SAML providers for ODBC/JDBC-based SAML authentication using the SAP HANA cockpit, SAP HANA studio or SQL, always use the SAP HANA XS Administration Tool to configure SAML providers that will be used for **HTTP access via the XS classic server**.

1. In the SAP HANA cockpit, navigate to the [Overview](#) page and choose the [Manage SAML providers](#) link.
2. Add a new identity provider.
 - a. Enter the name of the identity provider.

The following naming conventions apply:

 - Spaces and special characters except underscore (_) are not permitted.
 - The name must start with a letter.
 - The name cannot exceed 127 characters.
 - b. Enter the entity ID.
 - c. Select the appropriate X.509 certificate.

i Note

It is not possible to enter the issuer and subject distinguished names (DNs) manually. If the certificate is not available, click [Go to Certificate Store](#) and import it. Then, return to the [SAML Identity Provider](#) page and start again. For more information, see the section in importing a trusted certificate into the certificate store.

- d. Click [Add](#).

Results

The identity provider is now available for mapping to individual database users. You can do this when you create the database user. Alternatively, if the database user already exists, you can change their authentication details.

Related Information

[Import a Trusted Certificate into the Certificate Store \[page 909\]](#)

[Create a Database User \[page 792\]](#)

[Change a Database User \[page 802\]](#)

[Maintaining SAML Providers \(HTTP Access via XS Classic Server\) \[page 1559\]](#)

7.2.3.2.5 Configure the Trust Store for SAP Logon Tickets and Assertions

If you are integrating SAP HANA system into a landscape that uses SAP logon or assertion tickets for user authentication, you must configure SAP HANA to accept logon/assertion tickets.

Prerequisites

- If you are using certificate collections and certificates stored directly in the database (recommended), you have all the necessary privileges. For more information, see the section on managing client certificates.
- If you are using a trust store located in the file system, you have the system privilege INIFILE ADMIN.

Context

SAP HANA validates incoming logon/assertion tickets against certificates signed by a trusted Certification Authority (CA) stored in a dedicated trust store. This trust store must contain all root certificate(s) used to validate logon/assertion tickets. We recommend creating a certificate collection with the purpose **SAP LOGON** and the required certificates directly in the database.

It is also possible to use a trust store located in the file system. The default location of the trust store in the file system depends on the cryptographic library configured for SSL:

- `$SECUDIR/saplogon.pse` (CommonCryptoLib)

i Note

The saplogon.pse trust store is available automatically.

- `$HOME/.ssl/saplogon.pem` (OpenSSL)

i Note

Deprecated: OpenSSL is deprecated. You must migrate to CommonCryptoLib. For more information, see SAP Note 2093286.

If necessary, you can change the location of this trust store in the `indexserver.ini` system properties file.

⚠ Caution

By default, the same trust store in the file system is shared by all databases. Different PSEs must be explicitly configured for tenant databases.

Procedure

Configure the trust store:

Option	Description
In-database certificate collection	In the database, create a certificate collection with the purpose SAP LOGON . For more information, see the section on managing client certificates in the SAP HANA database in the <i>SAP HANA Administration Guide</i> .
File system based	<ol style="list-style-type: none">1. In the <code>indexserver.ini</code> file, change the value of the <code>[authentication] saplogontickettruststore</code> parameter.2. Restart the system.

Related Information

[Managing Client Certificates \[page 900\]](#)

[SAP Note 2093286](#)

7.2.3.2.6 Configure an LDAP Server Connection for LDAP User Authentication

To enable LDAP user authentication, you set up a connection to an LDAP server by creating an LDAP provider in the SAP HANA database. Depending on your requirements, you configure the LDAP server to authenticate users only, or to authenticate and authorize users. For LDAP-authenticated users, you can also enable the automatic creation of users in SAP HANA.

Prerequisites

- An LDAP v3 compliant server
- You have the system privilege `LDAP ADMIN`.
- A certificate collection with purpose `LDAP` exists in the database and the certificate of the Certificate Authority (CA) that signed the certificate used by the LDAP server has been added. This is required to enable secure communication between SAP HANA and the LDAP server using the TLS/SSL protocol.

⚠ Caution

You must secure communication to protect the transmission of user passwords between SAP HANA and the LDAP server. You must also secure communication between clients and SAP HANA to protect data transmitted between the client and SAP HANA, as well as to ensure that the client is connecting to the expected SAP HANA (server authentication).

- LDAP authentication is an active authentication mechanism in the SAP HANA database. You can verify this by checking the value of the parameter `[authentication] authentication_methods` in the `global.ini` configuration file.
- LDAP groups have been mapped to roles in SAP HANA. For more information, see the section on configuring LDAP group authorization.

Procedure

1. Create and configure the LDAP provider in the SAP HANA database using the `CREATE LDAP PROVIDER` statement.

For more information, see the section on LDAP provider details.

→ Remember

You can change the configuration of an LDAP provider using the `ALTER LDAP PROVIDER`.

2. Verify the configuration of the LDAP provider.

You do this using the `VALIDATE LDAP PROVIDER` statement. This will help you to identify and rectify any missing or incorrect settings.

🔗 Example

Example 1:

This example verifies that based on the current LDAP provider configuration and role-to-group mappings in SAP HANA, the SAP HANA user `john` will be granted SAP HANA roles once successfully authenticated:

```
VALIDATE LDAP PROVIDER my_ldap_provider CHECK USER 'john'
```

Example 2:

This example verifies that based on the current LDAP provider configuration and role-to-group mappings in SAP HANA, the user `julie` will be created in the SAP HANA database and granted SAP HANA roles once successfully authenticated (using LDAP authentication).

```
VALIDATE LDAP PROVIDER my_ldap_provider CHECK USER CREATION FOR LDAP USER 'julie'
```

Next Steps

If you want authenticated users to be authorized on the basis of their LDAP group membership and you did **not** configure the LDAP provider for automatic user creation, configure SAP HANA users for LDAP group authorization. For more information, see the section on configuring LDAP group authorization.

Related Information

[LDAP Provider Details \[page 821\]](#)

[Configure LDAP Group Authorization \[page 819\]](#)

[Managing Client Certificates \[page 900\]](#)

7.2.3.2.6.1 Securing Communication Between the SAP HANA Server and LDAP Server Using SSL/TLS

Use SSL/TLS to secure the communication between the SAP HANA Server and LDAP Server. The two methods to encrypt LDAP connection with SSL/TLS are:

- LDAP over SSL: Traditionally, secured LDAP connections were handled on a separate port (port 636) and required the use of `ldaps://`. In this approach, SSL/TLS handshake is completed before any LDAP protocol messages are exchanged.
To use LDAP over SSL, the URLs specified in the LDAP provider must start with `ldaps://` and SSL clause must be set to OFF.
- LDAP with STARTTLS: In this method, the LDAP connection is initiated as unsecured; however, STARTTLS upgrades the connection to encrypted during the connection. It allows LDAP server to handle secured and unsecured connections using the same port (port 389). Note that the connection is secured before user credentials are transmitted. To use STARTTLS, the URLs specified in the LDAP provider must start with `ldap://` and SSL clause must be set to ON.

7.2.4 User Authorization

After successful logon, the user's authorization to perform the requested operations on the requested objects is verified.

To perform operations in the SAP HANA database, a database user must have the necessary privileges. Users must have both the privilege(s) to perform the operation and to access the resources (such as schemas and tables) to which the operation applies. Privileges can be granted to database users either directly, or indirectly through roles that they have been granted. In this case, the privileges are inherited. Roles are the standard mechanism of granting privileges to users.

i Note

For some administration tasks, such as start-up, the credentials of the SAP operating system user (<sid>adm) are also required.

[Privileges \[page 740\]](#)

Several privilege types are used in SAP HANA (system, object, analytic, package, and application).

[System Privileges \[page 742\]](#)

System privileges control general system activities.

[Object Privileges \[page 748\]](#)

Object privileges are SQL privileges that are used to allow access to and modification of database objects.

[Analytic Privileges \[page 754\]](#)

Analytic privileges grant different users access to different portions of data in the same view based on their business role. Within the definition of an analytic privilege, the conditions that control which data users see is either contained in an XML document or defined using SQL.

[Package Privileges \[page 756\]](#)

Package privileges authorize actions on individual packages in the classic SAP HANA repository.

[Application Privileges \[page 758\]](#)

In SAP HANA XS classic, application privileges define the authorization level required for access to an SAP HANA XS classic application, for example, to start the application or view particular functions and screens.

[Database Roles \[page 759\]](#)

A database role is a collection of privileges that can be granted to either a database user or another role in runtime.

[System Views for Verifying Users' Authorization \[page 764\]](#)

You can query several system views to get detailed information about exactly which privileges and roles users have and how they come to have them. This can help you to understand why a user is authorized to perform particular actions, access particular data, or not.

[Restrict Use of the CLIENT User Parameter \[page 767\]](#)

Allow only authorized technical users to overwrite the value of the `CLIENT` parameter for a database connection or the value of the `$$client$$` parameter in an SQL query.

[Resolve Errors Using the Authorization Dependency Viewer \[page 768\]](#)

You can use the authorization dependency viewer in the SAP HANA studio as a first step in troubleshooting authorization errors and invalid object errors for stored procedures and calculation views with complex dependency structures.

7.2.4.1 Privileges

Several privilege types are used in SAP HANA (system, object, analytic, package, and application).

Privilege Type	Applicable To	Target User	Description
System privilege	System, database	Administrators, developers	<p>System privileges control general system activities. They are mainly used for administrative purposes, such as creating schemas, creating and changing users and roles, performing data backups, managing licenses, and so on.</p> <p>System privileges are also used to authorize basic repository operations.</p> <p>System privileges granted to users in a particular tenant database authorize operations in that database only. The only exception is the system privileges DATABASE ADMIN, DATABASE STOP, and DATABASE START. These system privileges can only be granted to users of the system database. They authorize the execution of operations on individual tenant databases. For example, a user with DATABASE ADMIN can create and drop tenant databases, change the database-specific properties in configuration (*.ini) files, and perform database-specific backups.</p>
Object privilege	Database objects (schemas, tables, views, procedures and so on)	End users, technical users	<p>Object privileges are used to allow access to and modification of database objects, such as tables and views. Depending on the object type, different actions can be authorized (for example, SELECT, CREATE ANY, ALTER, DROP, and so on).</p> <p>Schema privileges are object privileges that are used to allow access to and modification of schemas and the objects that they contain.</p> <p>Source privileges are object privileges that are used to restrict access to and modification of remote data sources, which are connected through SAP HANA smart data access.</p> <p>Object privileges granted to users in a particular database authorize access to and modification of database objects in that database only. That is, unless cross-database access has been enabled for the user. This is made possible through the association of the requesting user with a remote identity on the remote database. For more information, see <i>Cross-Database Authorization in Tenant Databases</i> in the <i>SAP HANA Security Guide</i>.</p>

Privilege Type	Applicable To	Target User	Description
Analytic privilege	Analytic views	End users	<p>Analytic privileges are used to allow read access to data in SAP HANA information models (that is, analytic views, attribute views, and calculation views) depending on certain values or combinations of values. Analytic privileges are evaluated during query processing.</p> <p>Analytic privileges granted to users in a particular database authorize access to information models in that database only.</p>
Package privilege	Packages in the classic repository of the SAP HANA database	Application and content developers working in the classic SAP HANA repository	<p>Package privileges are used to allow access to and the ability to work in packages in the classic repository of the SAP HANA database.</p> <p>Packages contain design time versions of various objects, such as analytic views, attribute views, calculation views, and analytic privileges.</p> <p>Package privileges granted to users in a particular database authorize access to and the ability to work in packages in the repository of that database only.</p> <div data-bbox="927 1048 1396 1305" style="background-color: #f0f0f0; padding: 10px;"> <p>i Note</p> <p>With SAP HANA XS advanced, source code and web content are not versioned and stored in the SAP HANA database, so package privileges are not used in this context. For more information, see <i>Authorization in SAP HANA XS Advanced</i>.</p> </div>
Application privilege	SAP HANA XS classic applications	Application end users, technical users (for SQL connection configurations)	<p>Developers of SAP HANA XS classic applications can create application privileges to authorize user and client access to their application. They apply in addition to other privileges, for example, object privileges on tables.</p> <p>Application privileges can be granted directly to users or roles in runtime in the SAP HANA studio. However, it is recommended that you grant application privileges to roles created in the repository in design time.</p> <div data-bbox="927 1686 1396 1975" style="background-color: #f0f0f0; padding: 10px;"> <p>i Note</p> <p>With SAP HANA XS advanced, application privileges are not used. Application-level authorization is implemented using OAuth and authorization scopes and attributes. For more information, see <i>Authorization in SAP HANA XS Advanced</i>.</p> </div>

i Note

An additional privilege type, privileges on users, can be granted to users. Privileges on users are SQL privileges that users can grant on their user. ATTACH DEBUGGER is the only privilege that can be granted on a user.

For example, User A can grant User B the privilege ATTACH DEBUGGER to allow User B debug SQLScript code in User A's session. User A is only user who can grant this privilege. Note that User B also needs the object privilege DEBUG on the relevant SQLScript procedure.

For more information, see the section on debugging procedures in the *SAP HANA Developer Guide*.

7.2.4.2 System Privileges

System privileges control general system activities.

System privileges are mainly used to authorize users to perform administrative actions, including:

- Creating and deleting schemas
- Managing users and roles
- Performing data backups
- Monitoring and tracing
- Managing licenses

System privileges are also used to authorize basic repository operations, for example:

- Importing and exporting content
- Maintaining delivery units (DU)

System privileges granted to users in a particular database authorize operations in that database only. The only exception is the system privileges DATABASE ADMIN, DATABASE STOP, and DATABASE START . These system privileges can only be granted to users of the system database. They authorize the execution of operations on individual tenant databases. For example, a user with DATABASE ADMIN can create and drop tenant databases, change the database-specific properties in configuration (*.ini) files, and perform database-specific or full-system data backups.

Related Information

[System Privileges \(Reference\) \[page 743\]](#)

7.2.4.2.1 System Privileges (Reference)

System privileges control general system activities.

General System Privileges

System privileges restrict administrative tasks. The following table describes the supported system privileges in an SAP HANA database.

System Privilege	Description
ADAPTER ADMIN	Controls the execution of the following adapter-related statements: CREATE ADAPTER, DROP ADAPTER, and ALTER ADAPTER. It also allows access to the ADAPTERS and ADAPTER_LOCATIONS system views.
AGENT ADMIN	Controls the execution of the following agent-related statements: CREATE AGENT, DROP AGENT, and ALTER AGENT. It also allows access to the AGENTS and ADAPTER_LOCATIONS system views.
ATTACH DEBUGGER	Authorizes debugging across different user sessions. For example, userA can grant ATTACH DEBUGGER to userB to allow userB to debug a procedure in userA's session (userB still needs DEBUG privilege on the procedure, however).
AUDIT ADMIN	Controls the execution of the following auditing-related statements: CREATE AUDIT POLICY, DROP AUDIT POLICY, and ALTER AUDIT POLICY, as well as changes to the auditing configuration. It also allows access to the AUDIT_LOG, XSA_AUDIT_LOG, and ALL_AUDIT_LOG system views.
AUDIT OPERATOR	Authorizes the execution of the following statement: ALTER SYSTEM CLEAR AUDIT LOG. It also allows access to the AUDIT_LOG system view.
BACKUP ADMIN	Authorizes BACKUP and RECOVERY statements for defining and initiating backup and recovery procedures. It also authorizes changing system configuration options with respect to backup and recovery.
BACKUP OPERATOR	Authorizes the BACKUP statement to initiate a backup.

System Privilege	Description
CATALOG READ	Authorizes unfiltered access to the data in the system views that a user has already been granted the SELECT privilege on. Normally, the content of these views is filtered based on the privileges of the user. CATALOG READ does not allow a user to view system views on which they have not been granted the SELECT privilege.
CERTIFICATE ADMIN	Authorizes the changing of certificates and certificate collections that are stored in the database.
CLIENT PARAMETER ADMIN	Authorizes a user to override the value of the CLIENT parameter for a database connection or to overwrite the value of the \$\$client\$\$ parameter in a SQL query.
CREATE CLIENTSIDE ENCRYPTION KEYPAIR	Authorizes a user to create client-side encryption key pairs.
CREATE R SCRIPT	Authorizes the creation of a procedure by using the language R.
CREATE REMOTE SOURCE	Authorizes the creation of remote data sources by using the CREATE REMOTE SOURCE statement.
CREATE SCENARIO	Controls the creation of calculation scenarios and cubes (calculation database).
CREATE SCHEMA	Authorizes the creation of database schemas using the CREATE SCHEMA statement.
CREATE STRUCTURED PRIVILEGE	Authorizes the creation of structured (analytic privileges). Only the owner of the privilege can further grant or revoke that privilege to other users or roles.
CREDENTIAL ADMIN	Authorizes the use of the statements CREATE CREDENTIAL, ALTER CREDENTIAL, and DROP CREDENTIAL.
DATA ADMIN	Authorizes reading all data in the system views. It also enables execution of Data Definition Language (DDL) statements in the SAP HANA database. A user with this privilege cannot select or change data in stored tables for which they do not have access privileges, but they can drop tables or modify table definitions.
DATABASE ADMIN	Authorizes all statements related to tenant databases, such as CREATE, DROP, ALTER, RENAME, BACKUP, and RECOVERY.
DATABASE START	Authorizes a user to start any database in the system and to select from the M_DATABASES view.

System Privilege	Description
DATABASE STOP	Authorizes a user to stop any database in the system and to select from the M_DATABASES view.
DROP CLIENTSIDE ENCRYPTION KEYPAIR	Authorizes a user to drop other users' client-side encryption key pairs.
ENCRYPTION ROOT KEY ADMIN	Authorizes all statements related to management of root keys: Allows access to the system views pertaining to encryption (for example, ENCRYPTION_ROOT_KEYS, M_ENCRYPTION_OVERVIEW, M_PERSISTENCE_ENCRYPTION_STATUS, M_PERSISTENCE_ENCRYPTION_KEYS, and so on).
EXPORT	Authorizes EXPORT to a file on the SAP HANA server. The user must also have the SELECT privilege on the source tables to be exported.
EXTENDED STORAGE ADMIN	Authorizes the management of SAP HANA dynamic tiering and the creation of extended storage.
IMPORT	Authorizes the import activity in the database using the IMPORT statements. The user must also have the INSERT privilege on the target tables to be imported.
INIFILE ADMIN	Authorizes making changes to system settings.
LDAP ADMIN	Authorizes the use of the CREATE ALTER DROP VALIDATE LDAP PROVIDER statements.
LICENSE ADMIN	Authorizes the use of the SET SYSTEM LICENSE statement to install a new license.
LOG ADMIN	Authorizes the use of the ALTER SYSTEM LOGGING [ON OFF] statements to enable or disable the log flush mechanism.
MONITOR ADMIN	Authorizes the use of the ALTER SYSTEM statements for events.
OPTIMIZER ADMIN	Authorizes the use of the ALTER SYSTEM statements concerning SQL PLAN CACHE and ALTER SYSTEM UPDATE STATISTICS statements, which influence the behavior of the query optimizer.

System Privilege	Description
RESOURCE ADMIN	Authorizes statements concerning system resources (for example, the ALTER SYSTEM RECLAIM DATAVOLUME and ALTER SYSTEM RESET MONITORING VIEW statements). It also authorizes many of the statements available in the Management Console.
ROLE ADMIN	<p>Authorizes the creation and deletion of roles by using the CREATE ROLE and DROP ROLE statements. It also authorizes the granting and revoking of roles by using the GRANT and REVOKE statements.</p> <p>Activated repository roles, meaning roles whose creator is the predefined user _SYS_REPO, can neither be granted to other roles or users nor dropped directly. Not even users with the ROLE ADMIN privilege can do so. Check the documentation concerning activated objects.</p>
SAVEPOINT ADMIN	Authorizes the execution of a savepoint using the ALTER SYSTEM SAVEPOINT statement.
SCENARIO ADMIN	Authorizes all calculation scenario-related activities (including creation).
SERVICE ADMIN	Authorizes the ALTER SYSTEM [START CANCEL RECONFIGURE] statements for administering system services of the database.
SESSION ADMIN	Authorizes the ALTER SYSTEM commands concerning sessions to stop or disconnect a user session or to change session variables.
SSL ADMIN	Authorizes the use of the SET...PURPOSE SSL statement. It also allows access to the PSES system view.
STRUCTUREDPRIVILEGE ADMIN	Authorizes the creation, reactivation, and dropping of structured (analytic) privileges.
TENANT ADMIN	Authorizes the tenant operations performed by the ALTER SYSTEM [RESUME SUSPEND] TENANT statements.
TABLE ADMIN	Authorizes LOAD, UNLOAD and MERGE of tables and table placement.
TRACE ADMIN	Authorizes the use of the ALTER SYSTEM...TRACES statements for operations on database trace files and authorizes changing trace system settings.
TRUST ADMIN	Authorizes the use of statements to update the trust store.

System Privilege	Description
USER ADMIN	Authorizes the creation and modification of users by using the CREATE ALTER DROP USER statements.
VERSION ADMIN	Authorizes the use of the ALTER SYSTEM RECLAIM VERSION SPACE statement of the multi-version concurrency control (MVCC) feature.
WORKLOAD ADMIN	Authorizes execution of the workload class and mapping statements (for example, CREATE ALTER DROP WORKLOAD CLASS, and CREATE ALTER DROP WORKLOAD MAPPING).
WORKLOAD ANALYZE ADMIN	Used by the Analyze Workload, Capture Workload, and Replay Workload applications when performing workload analysis.
WORKLOAD CAPTURE ADMIN	Authorizes access to the monitoring view M_WORKLOAD_CAPTURES to see the current status of capturing and captured workloads, as well of execution of actions with the WORKLOAD_CAPTURE procedure.
WORKLOAD REPLAY ADMIN	Authorizes access to the monitoring views M_WORKLOAD_REPLAY_PREPROCESSES and M_WORKLOAD_REPLAYS to see current status of preprocessing, preprocessed, replaying, and replayed workloads, as well as the execution of actions with the WORKLOAD_REPLAY procedure.
<identifier>.<identifier>	Components of the SAP HANA database can create new system privileges. These privileges use the component-name as the first identifier of the system privilege and the component-privilege-name as the second identifier.

Repository System Privileges

i Note

The following privileges authorize actions on individual packages in the SAP HANA repository, used in the SAP HANA Extended Services (SAP HANA XS) classic development model. With SAP HANA XS advanced, source code and web content are no longer versioned and stored in the repository of the SAP HANA database.

System Privilege	Description
REPO.EXPORT	Authorizes the export of delivery units for example

System Privilege	Description
REPO.IMPORT	Authorizes the import of transport archives
REPO.MAINTAIN_DELIVERY_UNITS	Authorizes the maintenance of delivery units (DU, DU vendor and system vendor must be the same)
REPO.WORK_IN_FOREIGN_WORKSPACE	Authorizes work in a foreign inactive workspace
REPO.CONFIGURE	Authorize work with SAP HANA Change Recording, which is part of SAP HANA Application Lifecycle Management
REPO.MODIFY_CHANGE	
REPO.MODIFY_OWN_CONTRIBUTION	
REPO.MODIFY_FOREIGN_CONTRIBUTION	

7.2.4.3 Object Privileges

Object privileges are SQL privileges that are used to allow access to and modification of database objects.

For each SQL statement type (for example, SELECT, UPDATE, or CALL), a corresponding object privilege exists. If a user wants to execute a particular statement on a simple database object (for example, a table), he or she must have the corresponding object privilege for either the actual object itself, or the schema in which the object is located. This is because the schema is an object type that contains other objects. A user who has object privileges for a schema automatically has the same privileges for all objects currently in the schema and any objects created there in the future.

Object privileges are not only grantable for database catalog objects such as tables, views and procedures. Object privileges can also be granted for non-catalog objects such as development objects in the repository of the SAP HANA database.

Initially, the owner of an object and the owner of the schema in which the object is located are the only users who can access the object and grant object privileges on it to other users.

An object can therefore be accessed only by the following users:

- The owner of the object
- The owner of the schema in which the object is located
- Users to whom the owner of the object has granted privileges
- Users to whom the owner of the parent schema has granted privileges

Caution

The database owner concept stipulates that when a database user is deleted, all objects created by that user and privileges granted to others by that user are also deleted. If the owner of a schema is deleted, all objects in the schema are also deleted even if they are owned by a different user. All privileges on these objects are also deleted.

Note

The owner of a table can change its ownership with the `ALTER TABLE` SQL statement. In this case, the new owner becomes the grantor of all privileges on the table granted by the original owner. The original owner is

also automatically granted all privileges for the table with the new owner as grantor. This ensures that the original owner can continue to work with the table as before.

Authorization Check on Objects with Dependencies

The authorization check for objects defined on other objects (that is, stored procedures and views) is more complex. In order to be able to access an object with dependencies, both of the following conditions must be met:

- The user trying to access the object must have the relevant object privilege on the object as described above.
- The user who created the object must have the required privilege on all underlying objects **and** be authorized to grant this privilege to others.

If this second condition is not met, only the owner of the object can access it. He cannot grant privileges on it to any other user. This cannot be circumvented by granting privileges on the parent schema instead. Even if a user has privileges on the schema, he will still not be able to access the object.

i Note

This applies to procedures created in DEFINER mode only. This means that the authorization check is run against the privileges of the user who created the object, not the user accessing the object. For procedures created in INVOKER mode, the authorization check is run against the privileges of the accessing user. In this case, the user must have privileges not only on the object itself but on all objects that it uses.

→ Tip

The SAP HANA studio provides a graphical feature, the authorization dependency viewer, to help troubleshoot authorization errors for object types that typically have complex dependency structures: stored procedures and calculation views.

Related Information

[Resolve Errors Using the Authorization Dependency Viewer \[page 768\]](#)

[Object Privileges \(Reference\) \[page 749\]](#)

7.2.4.3.1 Object Privileges (Reference)

Object privileges are used to allow access to and modification of database objects, such as tables and views.

The following table describes the supported object privileges in an SAP HANA database.

Object Privilege	Command Types	Applies to	Privilege Description
ALL PRIVILEGES	DDL & DML	<ul style="list-style-type: none"> • Schemas • Tables • Views 	<p>This privilege is a collection of all Data Definition Language (DDL) and Data Manipulation Language (DML) privileges that the grantor currently possesses and is allowed to grant further. The privilege it grants is specific to the particular object being acted upon.</p> <p>This privilege collection is dynamically evaluated for the given grantor and object.</p>
ALTER	DDL	<ul style="list-style-type: none"> • Schemas • Tables • Views • Functions/procedures 	Authorizes the ALTER statement for the object.
CREATE ANY	DDL	<ul style="list-style-type: none"> • Schemas • Tables • Views • Sequences • Functions/procedures • Remote sources • Graph workspaces 	Authorizes all CREATE statements for the object.
CREATE VIRTUAL FUNCTION	DDL	<ul style="list-style-type: none"> • Remote sources 	Authorizes creation of virtual functions (the REFERENCES privilege is also required).
CREATE VIRTUAL PROCEDURE	DDL	<ul style="list-style-type: none"> • Remote sources 	Authorizes creation of virtual procedure to create and run procedures on a remote source.
CREATE VIRTUAL PACKAGE	DDL	<ul style="list-style-type: none"> • Schemas 	Authorizes creation of virtual packages that can be run on remote sources.
CREATE VIRTUAL TABLE	DDL	<ul style="list-style-type: none"> • Remote sources 	Authorizes the creation of proxy tables pointing to remote tables from the source entry.

Object Privilege	Command Types	Applies to	Privilege Description
CREATE TEMPORARY TABLE	DDL	<ul style="list-style-type: none"> Schemas 	Authorizes the creation of a temporary local table, which can be used as input for procedures, even if the user does not have the CREATE ANY privilege for the schema.
DEBUG	DML	<ul style="list-style-type: none"> Schemas Calculation Views Functions/procedures 	Authorizes debug functionality for the procedure or calculation view or for the procedures and calculation views of a schema.
DEBUG MODIFY	DDL	<ul style="list-style-type: none"> Functions/procedures 	For internal use only.
DELETE	DML	<ul style="list-style-type: none"> Schemas Tables Views Functions/procedures 	<p>Authorizes the DELETE and TRUNCATE statements for the object.</p> <p>While DELETE applies to views, it only applies to updatable views (that is, views that do not use a join, do not contain a UNION, and do not use aggregation).</p>
DROP	DDL	<ul style="list-style-type: none"> Schemas Tables Views Sequences Functions/procedures Remote sources Graph workspaces 	Authorizes the DROP statements for the object.
EXECUTE	DML	<ul style="list-style-type: none"> Schemas Functions/procedures 	Authorizes the execution of a SQLScript function or a database procedure by using the CALLS or CALL statement respectively. It also allows a user to execute a virtual function.
INDEX	DDL	<ul style="list-style-type: none"> Schemas Tables 	Authorizes the creation, modification, or dropping of indexes for the object.

Object Privilege	Command Types	Applies to	Privilege Description
INSERT	DML	<ul style="list-style-type: none"> • Schemas • Tables • Views 	<p>Authorizes the INSERT statement for the object.</p> <p>The INSERT and UPDATE privilege are both required on the object to allow the REPLACE and UPSERT statements to be used.</p> <p>While INSERT applies to views, it only applies to updatable views (views that do not use a join, do not contain a UNION, and do not use aggregation).</p>
REFERENCES	DDL	<ul style="list-style-type: none"> • Schemas • Tables 	<p>Authorizes the usage of all tables in this schema or this table in a foreign key definition, or the usage of a personal security environment (PSE). It also allows a user to reference a virtual function package.</p>
SELECT	DML	<ul style="list-style-type: none"> • Schemas • Tables • Views • Sequences • Graph workspaces 	<p>Authorizes the SELECT statement for the object or the usage of a sequence.</p> <p>When selection from system-versioned tables, users must have SELECT on both the table and its associated history table.</p>
SELECT CDS METADATA	DML	<ul style="list-style-type: none"> • Schemas • Tables 	<p>Authorizes access to CDS metadata from the catalog.</p>
SELECT METADATA	DML	<ul style="list-style-type: none"> • Schemas • Tables 	<p>Authorizes access to the complete metadata of all objects in a schema (including procedure and view definitions), including objects that may be located in other schemas.</p>

Object Privilege	Command Types	Applies to	Privilege Description
TRIGGER	DDL	<ul style="list-style-type: none"> Schemas Tables 	Authorizes the CREATE TRIGGER/DROP TRIGGER statement for the specified table or the tables in the specified schema.
UNMASKED	DML	<ul style="list-style-type: none"> Schemas Views Tables 	Authorizes access to masked data in user-defined views and tables. This privilege is required to view the original data in views and tables that are defined by using the WITH MASK clause.
UPDATE	DML	<ul style="list-style-type: none"> Schemas Tables Views 	While UPDATE applies to views, it only applies to updatable views (views that do not use a join, do not contain a UNION, and do not use aggregation).
USERGROUP OPERATOR	DML	<ul style="list-style-type: none"> User groups 	<p>Authorizes a user to change the settings for a user group, and to add and remove users to/from a user group.</p> <p>Users with the USERGROUP OPERATOR privilege can also create and drop users, but only within the user group they have the USERGROUP OPERATOR privilege on (CREATE USER <user_name> SET USERGROUP <usergroup_name>).</p> <p>A user can have the USERGROUP OPERATOR privilege on more than one user group, and a user group can have more than one user with the USERGROUP OPERATOR privilege on it.</p>

Object Privilege	Command Types	Applies to	Privilege Description
<code><identifier>.<identifier></code>	DDL		Components of the SAP HANA database can create new object privileges. These privileges use the component-name as first identifier of the system privilege and the component-privilege-name as the second identifier.

7.2.4.4 Analytic Privileges

Analytic privileges grant different users access to different portions of data in the same view based on their business role. Within the definition of an analytic privilege, the conditions that control which data users see is either contained in an XML document or defined using SQL.

Standard object privileges (`SELECT`, `ALTER`, `DROP`, and so on) implement coarse-grained authorization at object level only. Users either have access to an object, such as a table, view or procedure, or they don't. While this is often sufficient, there are cases when access to data in an object depends on certain values or combinations of values. Analytic privileges are used in the SAP HANA database to provide such fine-grained control at row level of which data individual users can see within the same view.

❁ Example

Sales data for all regions are contained within one analytic view. However, regional sales managers should only see the data for their region. In this case, an analytic privilege could be modeled so that they can all query the view, but only the data that each user is authorized to see is returned.

Creation of Analytic Privileges

Although analytic privileges can be created directly as catalog objects in runtime, we recommend creating them as design-time objects that become catalog objects on deployment (database artifact with file suffix `.hdbanalyticprivilege`). In an SAP HANA XS classic environment, analytic privileges are created in the built-in repository of the SAP HANA database using either the SAP HANA Web Workbench or the SAP HANA studio. In an SAP HANA XS advanced environment, they are created using the SAP Web IDE and deployed using SAP HANA deployment infrastructure (SAP HANA DI).

i Note

HDI supports only SQL-based analytic privileges (see below). Furthermore, due to the container-based model of HDI, where each container corresponds to a database schema, analytic privileges created in HDI are schema specific.

XML- Versus SQL-Based Analytic Privileges

Before you implement row-level authorization using analytic privileges, you need to decide which type of analytic privilege is suitable for your scenario. In general, SQL-based analytic privileges allow you to more easily formulate complex filter conditions using sub-queries that might be cumbersome to model using XML-based analytic privileges.

→ Recommendation

SAP recommends the use of SQL-based analytic privileges. Using the *SAP HANA Modeler* perspective of the SAP HANA studio, you can migrate XML-based analytic privileges to SQL-based analytic privileges. For more information, see the SAP HANA Modeling Guide (For SAP HANA Studio).

i Note

As objects created in the repository, XML-based analytic privileges are deprecated as of SAP HANA SPS 02. For more information, see SAP Note 2465027.

The following are the main differences between XML-based and SQL-based analytic privileges:

Feature	SQL-Based Analytic Privileges	XML-Based Analytic Privileges
Control of read-only access to SAP HANA information models: <ul style="list-style-type: none"> • Attribute views • Analytic views • Calculation views 	Yes	Yes
Control of read-only access to SQL views	Yes	No
Control of read-only access to database tables	No	No
Design-time modeling using the SAP HANA Web-based Workbench or the <i>SAP HANA Modeler</i> perspective of the SAP HANA studio	Yes	Yes
Design-time modeling using the SAP Web IDE for SAP HANA	Yes	No
Transportable	Yes	Yes
HDI support	Yes	No
Complex filtering	Yes	No

i Note

This corresponds to development in an SAP HANA XS classic environment using the SAP HANA repository.

i Note

This corresponds to development in an SAP HANA XS advanced environment using HDI.

Enabling an Authorization Check Based on Analytic Privileges

All column views modeled and activated in the SAP HANA modeler and the SAP HANA Web-based Development Workbench automatically enforce an authorization check based on analytic privileges. XML-based analytic privileges are selected by default, but you can switch to SQL-based analytic privileges.

Column views created using SQL must be explicitly registered for such a check by passing the relevant parameter:

- `REGISTERVIEWFORAPCHECK` for a check based on XML-based analytic privileges
- `STRUCTURED PRIVILEGE CHECK` for a check based on SQL-based analytic privileges

SQL views must always be explicitly registered for an authorization check based on analytic privileges by passing the `STRUCTURED PRIVILEGE CHECK` parameter.

i Note

It is not possible to enforce an authorization check on the same view using both XML-based and SQL-based analytic privileges. However, it is possible to build views with different authorization checks on each other.

Related Information

[SAP Note 2465027](#)

7.2.4.5 Package Privileges

Package privileges authorize actions on individual packages in the classic SAP HANA repository.

i Note

With SAP HANA XS advanced, source code and web content are not versioned and stored in the SAP HANA database, so package privileges are not used in this context.

Privileges granted on a repository package are implicitly assigned to the design-time objects in the package, as well as to all sub-packages. Users are only allowed to maintain objects in a repository package if they have the necessary privileges for the package in which they want to perform an operation, for example to read or write to an object in that package. To be able perform operations in all packages, a user must have privileges on the root package `.REPO_PACKAGE_ROOT`.

→ Recommendation

We recommend that package privileges be granted on a single package or a small number of specific packages belonging to your organization, rather than on the complete repository.

If the user authorization check establishes that a user does not have the necessary privileges to perform the requested operation in a specific package, the authorization check is repeated on the parent package and recursively up the package hierarchy to the root level of the repository. If the user does not have the necessary

privileges for any of the packages in the hierarchy chain, the authorization check fails and the user is not permitted to perform the requested operation.

In the context of repository package authorizations, there is a distinction between native packages and imported packages.

Privileges for Native Repository Packages

A native repository package is created in the current SAP HANA system and expected to be edited in the current system. To perform application-development tasks on **native** packages in the SAP HANA repository, developers typically need the privileges listed in the following table:

Package Privilege	Description
REPO.READ	Read access to the selected package and design-time objects (both native and imported)
REPO.EDIT_NATIVE_OBJECTS	Authorization to modify design-time objects in packages originating in the system the user is working in
REPO.ACTIVATE_NATIVE_OBJECTS	Authorization to activate/reactivate design-time objects in packages originating in the system the user is working in
REPO.MAINTAIN_NATIVE_PACKAGES	Authorization to update or delete native packages, or create sub-packages of packages originating in the system in which the user is working

Privileges for Imported Repository Packages

An imported repository package is created in a remote SAP HANA system and imported into the current system. To perform application-development tasks on **imported** packages in the SAP HANA repository, developers need the privileges listed in the following table:

Note

It is not recommended to work on imported packages. Imported packages should only be modified in exceptional cases, for example, to carry out emergency repairs.

Package Privilege	Description
REPO.READ	Read access to the selected package and design-time objects (both native and imported)
REPO.EDIT_IMPORTED_OBJECTS	Authorization to modify design-time objects in packages originating in a system other than the one in which the user is currently working
REPO.ACTIVATE_IMPORTED_OBJECTS	Authorization to activate (or reactivate) design-time objects in packages originating in a system other than the one in which the user is currently working

Package Privilege	Description
REPO.MAINTAIN_IMPORTED_PACKAGES	Authorization to update or delete packages, or create sub-packages of packages, which originated in a system other than the one in which the user is currently working

7.2.4.6 Application Privileges

In SAP HANA XS classic, application privileges define the authorization level required for access to an SAP HANA XS classic application, for example, to start the application or view particular functions and screens.

Note

With SAP HANA XS advanced, application privileges are not used. Application-level authorization is implemented using OAuth and authorization scopes and attributes.

Application privileges can be assigned to an individual user or to a group of users, for example, in a role. The role can also be used to assign system, object, package, and analytic privileges. You can use application privileges to provide different levels of access to the same application, for example, to provide advanced maintenance functions for administrators and view-only capabilities to normal users.

If you want to define application-specific privileges, you need to understand and maintain the relevant sections in the following design-time artifacts:

- Application-privileges file (.xsprivileges)
- Application-access file (.xsaccess)
- Role-definition file (<RoleName>.hdbrole)

Application privileges can be assigned to users individually or by means of a user **role**, for example, with the “*application privilege*” keyword in a role-definition file (<RoleName>.hdbrole) as illustrated in the following code. You store the roles as design-time artifacts within the application package structure they are intended for, for example, acme.com.hana.xs.appl.roles.

```
role acme.com.hana.xs.appl.roles::Display
{
  application privilege: acme.com.hana.xs.appl::Display;
  application privilege: acme.com.hana.xs.appl::View;
  catalog schema "ACME_XS_APP1": SELECT;
  package acme.com.hana.xs.appl: REPO.READ;
  package ".REPO_PACKAGE_ROOT" : REPO.READ;
  catalog sql object "_SYS_REPO"."PRODUCTS": SELECT;
  catalog sql object "_SYS_REPO"."PRODUCT_INSTANCES": SELECT;
  catalog sql object "_SYS_REPO"."DELIVERY_UNITS": SELECT;
  catalog sql object "_SYS_REPO"."PACKAGE_CATALOG": SELECT;
  catalog sql object "ACME_XS_APPL"."acme.com.hana.xs.appl.db::SYSTEM_STATE":
  SELECT, INSERT, UPDATE, DELETE;
}
```

The application privileges referenced in the role definition (for example, Display and View) are actually defined in an application-specific .xsprivileges file, as illustrated in the following example, which also contains entries for additional privileges that are not explained here.

i Note

The `.xsprivileges` file must reside in the package of the application to which the privileges apply.

The package where the `.xsprivileges` resides defines the scope of the application privileges; the privileges specified in the `.xsprivileges` file can only be used in the package where the `.xsprivileges` resides (or any sub-packages). This is checked during activation of the `.xsaccess` file and at runtime in the by the XS JavaScript API `$.session.(has|assert)AppPrivilege()`.

```
{
  "privileges" : [
    { "name" : "View", "description" : "View Product Details" },
    { "name" : "Configure", "description" : "Configure Product Details" },
    { "name" : "Display", "description" : "View Transport Details" },
    { "name" : "Administrator", "description" : "Configure/Run Everything" },
    { "name" : "ExecuteTransport", "description" : "Run Transports"},
    { "name" : "Transport", "description" : "Transports"}
  ]
}
```

The privileges are **authorized** for use with an application by inserting the *authorization* keyword into the corresponding `.xsaccess` file, as illustrated in the following example. Like the `.xsprivileges` file, the `.xsaccess` file must reside either in the root package of the application to which the privilege authorizations apply or the specific subpackage which requires the specified authorizations.

i Note

If a privilege is inserted into the `.xsaccess` file as an authorization requirement, a user must have this privilege to access the application package where the `.xsaccess` file resides. If there is more than one privilege, the user must have at least one of these privileges to access the content of the package.

```
{
  "prevent_xsrp": true,
  "exposed": true,
  "authentication": {
    "method": "Form"
  },
  "authorization": [
    "acme.com.hana.xs.appl:Display",
    "acme.com.hana.xs.appl:Transport"
  ]
}
```

7.2.4.7 Database Roles

A database role is a collection of privileges that can be granted to either a database user or another role in runtime.

A role typically contains the privileges required for a particular function or task, for example:

- Business end users reading reports using client tools such as Microsoft Excel
- Modelers creating models and reports
- Database administrators operating and maintaining the database and its users

Privileges can be granted directly to users of the SAP HANA database. However, roles are the standard mechanism of granting privileges as they allow you to implement complex, reusable authorization concepts that can be modeled on business roles.

Creation of Roles

Roles in the SAP HANA database can exist as runtime objects only (catalog roles), or as design-time objects that become catalog objects on deployment (database artifact with file suffix `.hdbrrole`).

In an SAP HANA XS classic environment, database roles are created in the built-in repository of the SAP HANA database using either the SAP HANA Web Workbench or the SAP HANA studio. These are also referred to as repository roles. In an SAP HANA XS advanced environment, design-time roles are created using the SAP Web IDE and deployed using SAP HANA deployment infrastructure (SAP HANA DI, or HDI).

i Note

Due to the container-based model of HDI where each container corresponds to a database schema, HDI roles, once deployed, are schema specific.

SAP HANA XS advanced has the additional concept of application roles and role collections. These are independent of database roles in SAP HANA itself. In the XS advanced context, SAP HANA database roles are used only to control access to database objects (for example, tables, views, and procedures) for XS advanced applications. For more information about the authorization concept of XS advanced, see the *SAP HANA Security Guide*.

Role Structure

A role can contain any number of the following privileges:

- **System privileges** for general system authorization, in particular administration activities
- **Object privileges** (for example, SELECT, INSERT, UPDATE) on database objects (for example, schemas, tables, views, procedures, and sequences)
- **Analytic privileges** on SAP HANA information models
- **Package privileges** on repository packages (for example, REPO.READ, REPO.EDIT_NATIVE_OBJECTS, REPO.ACTIVATE_NATIVE_OBJECTS)
- **Application privileges** for enabling access to SAP HANA-based applications developed in an SAP HANA XS classic environment

i Note

There are no HDI or XS advanced equivalents in the SAP HANA authorization concept for package privileges on repository packages and applications privileges on SAP HANA XS classic applications. For more information about the authorization concept of XS advanced, see the *SAP HANA Security Guide*.

A role can also contain other roles.

Roles Best Practices

For best performance of role operations, in particular, granting and revoking, keep the following basic rules in mind:

- Create roles with the smallest possible set of privileges for the smallest possible group of users who can share a role (principle of least privilege).
- Avoid granting object privileges at the schema level to a role if only a few objects in the schema are relevant for intended users.
- Avoid creating and maintaining all roles as a single user. Use several role administrator users instead.

7.2.4.7.1 Catalog Roles and Design-Time Roles Compared

It is possible to create roles as pure runtime objects that follow classic SQL principles or as design-time objects in either an SAP HANA XS advanced or classic environment.

In SAP HANA XS classic, database roles are created in the built-in repository of the SAP HANA database using either the SAP HANA Web Workbench or the SAP HANA studio. In SAP HANA XS advanced, design-time roles are created using the SAP Web IDE and deployed using SAP HANA deployment infrastructure (SAP HANA DI, or HDI).

i Note

SAP HANA XS classic and the SAP HANA repository are deprecated as of SAP HANA 2.0 SPS 02. For more information, see SAP Note 2465027.

The following table summarizes the differences between catalog roles and design-time roles:

Feature	Catalog Roles	Design-Time Repository Roles (XSC)	Design-Time HDI Roles (XSA)
Transportability	Roles cannot be transported between systems. They can only be created in runtime by users with the system privilege ROLE ADMIN.	Roles can be transported between systems using several transport options: <ul style="list-style-type: none">• SAP HANA Application Lifecycle Manager• The change and transport system (CTS+) of the SAP NetWeaver ABAP application server• SAP HANA Transport Container (HTC)	
Version management	No version management is possible.	The repository provides the basis for versioning. As repository objects, roles are stored in specific repository tables inside the database. This eliminates the need for an external version control system.	Roles are developed as design-time objects within a project stored in the GIT repository. The complete role history is therefore available in GIT.

Feature	Catalog Roles	Design-Time Repository Roles (XSC)	Design-Time HDI Roles (XSA)
Ownership	Roles are owned by the database user who creates them. If the creating user is dropped, any roles created in the user's own schema are also dropped.	The technical user <code>_SYS_REPO</code> is the owner of all roles created in the repository, not the database user who creates them. Therefore, roles are not directly associated with the creating user. To create a role, a database user needs only the privileges required to work in the repository.	<p>Roles are owned by the object owner of the container (technical user <code><container>#00</code>) where role development is taking place.</p> <p>Roles are not directly associated with the creating user. To create a role, a developer needs only the authorization required to work in the relevant space using the SAP Web IDE for SAP HANA.</p> <p>If roles for different purposes are developed in different containers, then roles are not all owned by the same technical user (as is the case with repository roles).</p>

Feature	Catalog Roles	Design-Time Repository Roles (XSC)	Design-Time HDI Roles (XSA)
Grant and revoke process	<p>Roles created in runtime are granted directly by the database user using the SQL GRANT and REVOKE statements.</p> <p>To grant privileges to a role, a user requires either the system privilege ROLE ADMIN, or all the privileges being granted to the role. In the latter case, if any of these privileges are revoked from the granting user, they are automatically revoked from the role.</p> <p>Roles can be revoked by the granting user or any user with the system privilege ROLE ADMIN.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>With the exception of roles granted by technical user <code>_SYS_REPO</code>, a user with ROLE ADMIN cannot revoke roles granted by technical users <code>SYS</code> and <code>_SYS*</code>.</p> </div> <p>If the granting user is dropped (not necessarily the role creator), all roles that he or she granted are not revoked.</p>	<p>Roles are granted and revoked using built-in procedures. Any administrator with the EXECUTE privilege on these can grant and revoke roles.</p>	<p>Roles are granted and revoked using database procedures provided in either the HDI container's or container group's API schema. Any container or container group administrator with the EXECUTE privilege on these procedures can grant and revoke roles.</p> <p>Any user with the system privilege ROLE ADMIN can also grant and revoke roles.</p>

In general, it is recommended that you model roles as design-time objects for the following reasons:

- Unlike roles created in runtime, roles created as design-time objects can be transported between systems. This is important for application development as it means that developers can model roles as part of their application's security concept and then ship these roles or role templates with the application. Being able to transport roles is also advantageous for modelers implementing complex access control on analytic content. They can model roles in a test system and then transport them into a production system. This avoids unnecessary duplication of effort.
- Roles created as design-time objects are not directly associated with a database user. They are created by technical users and granted through the execution of stored procedures. Any user with access to these procedures can grant and revoke a role. Roles created in runtime are granted directly by the database user and can be revoked only by the granting user or a user with the system privilege ROLE ADMIN. Additionally, if the database user is deleted, all roles that he or she granted are revoked. As database users correspond

to real people, this could impact the implementation of your authorization concept, for example, if an employee leaves the organization or is on vacation.

Catalog roles make sense in scenarios where user and role provisioning is carried out solely using a higher-level application that connects to SAP HANA through a technical user such as SAP Identity Management.

Related Information

[SAP Note 2465027](#)

7.2.4.8 System Views for Verifying Users' Authorization

You can query several system views to get detailed information about exactly which privileges and roles users have and how they come to have them. This can help you to understand why a user is authorized to perform particular actions, access particular data, or not.

→ Remember

You must have the system privilege CATALOG READ to query the following views.

System View	Query	Result
ACCESSIBLE_VIEWS	<pre>SELECT * from "PUBLIC"."ACCESSIBLE_VIEWS" where USER_NAME = '<user_name>';</pre>	All views that the user is authorized to access are returned.
EFFECTIVE_APPLICATION_PRIVILEGES	<pre>select * from "SYS"."EFFECTIVE_APPLICATION_PRIVILEGES" where USER_NAME='<user_name>;</pre>	All application privileges granted to the specified user both directly and indirectly through roles are returned separately.
EFFECTIVE_MASK_EXPRESSIONS	<pre>SELECT * FROM EFFECTIVE_MASK_EXPRESSIONS where ROOT_SCHEMA_NAME = '<schema_name>' and ROOT_OBJECT_NAME = '<object_name>' and ROOT_COLUMN_NAME = '<column_name>' and USER_NAME ='<user_name>;</pre>	All masked columns that the specified user can see in the specified view and the corresponding mask expressions

System View	Query	Result
EFFECTIVE_ROLES	<pre>SELECT * FROM "PUBLIC"."EFFECTIVE_ROLES" where USER_NAME = '<user_name>' AND ROLE_SCHEMA_NAME = '<schema_name of role>';</pre>	All roles granted to the specified user both directly and indirectly through other roles are returned separately.
	<p>i Note</p> <p>Schema name is optional.</p>	
EFFECTIVE_STRUCTURED_PRIVILEGES	<pre>SELECT * from "PUBLIC"."EFFECTIVE_STRUCTURED_PRIVILEGES" where ROOT_SCHEMA_NAME = '<schema>' AND ROOT_OBJECT_NAME = '<object_name>' AND USER_NAME = '<user_name>'</pre>	The analytic privileges that are applicable to the specified view are returned, including dynamic filter conditions if relevant. It is also indicated whether or not the specified user is authorized to access the view.
GRANTED_PRIVILEGES	<pre>SELECT * FROM "PUBLIC"."GRANTED_PRIVILEGES" where GRANTEE = '<user_name>';</pre>	Privileges granted directly to the specified user (or role) are returned. Privileges contained within granted roles are not shown.
	<p>i Note</p> <p>It is possible to query the privileges directly granted to a role by replacing where GRANTEE = '<USER>' with where GRANTEE = '<ROLE>'</p>	
GRANTED_ROLES	<pre>SELECT * FROM "PUBLIC"."GRANTED_ROLES" where GRANTEE = '<user/role_name>';</pre>	All roles granted directly to the specified user (or role) are returned. Roles contained within granted roles are not shown.
	<p>i Note</p> <p>It is possible to query the roles directly granted to a role by replacing where GRANTEE = '<USER>' with where GRANTEE = '<ROLE>'</p>	

7.2.4.9 Restrict Use of the CLIENT User Parameter

Allow only authorized technical users to overwrite the value of the `CLIENT` parameter for a database connection or the value of the `$$client$$` parameter in an SQL query.

Context

The `CLIENT` user parameter can be used to authorize named users in SAP HANA database. Only a user with the `USER ADMIN` system privilege can change the value of the `CLIENT` parameter already assigned to other users. However, at runtime, any user can assign an arbitrary value to the `CLIENT` parameter either by setting the corresponding session variable or passing the parameter via placeholder in a query.

While this is the desired behavior for technical users that work with multiple clients such as SAP Business Warehouse, S/4 HANA, or SAP Business Suite, it is problematic in named user scenarios if the `CLIENT` parameter is used to authorize access to data and not only to perform data filtering.

Procedure

1. Grant the system privilege `CLIENT PARAMETER ADMIN` to database users or roles who are permitted to access to the `CLIENT` user parameter (for example, technical users).

Sample Code

```
GRANT CLIENT PARAMETER ADMIN TO <user or role>;
```

2. In the `global.ini` configuration file, see the value of the `[authorization]` `secure_client_parameter` to `true`.

Results

Only users with the system privilege `CLIENT PARAMETER ADMIN` can overwrite the value of the `CLIENT` parameter for a database connection or the value of the `$$client$$` parameter in an SQL query.

Related Information

[SAP Note 2582162](#)

7.2.4.10 Resolve Errors Using the Authorization Dependency Viewer

You can use the authorization dependency viewer in the SAP HANA studio as a first step in troubleshooting authorization errors and invalid object errors for stored procedures and calculation views with complex dependency structures.

Prerequisites

You have the system privilege CATALOG READ.

Context

The authorization dependency viewer is a graphical tool that depicts the object dependency structure of stored procedures and calculation views together with the SQL authorization status of the object owner along the dependency paths.

You can use the authorization dependency viewer as a first step in troubleshooting the following authorization errors and invalid object errors for these object types:

- NOT AUTHORIZED (258)
- INVALIDATED VIEW (391)
- INVALIDATED PROCEDURE (430)

Authorization or invalid object errors occur if the object owner does not have all the required privileges on all underlying objects on which the object depends (for example, tables, views, and procedures). The object owner must have both the appropriate SQL object privilege (for example, EXECUTE, SELECT) and the authorization to grant the object privilege to others (that is, WITH GRANT OPTION is set).

The authorization dependency viewer helps you to identify where there are invalid authorization dependencies in the object structure. This is particularly useful for objects with large and complex dependency structures.

→ Recommendation

Use the authorization dependency viewer only with procedures with security mode DEFINER. Procedures with security mode INVOKER are not validated correctly.

⚠ Caution

The authorization dependency viewer simply shows you which privileges are missing. Grant missing privileges with due care.

Procedure

1. Open the procedure or calculation view in the authorization dependency viewer:
 - a. Navigate to the object in the *Systems* view.
 - b. In the context menu, choose *Show Authorization*.

The object dependency structure is displayed as a hierarchical tree. Each node in the structure represents a database object. The same database object may appear multiple times if it is referenced at different levels of the tree. The lines connecting the nodes indicate the nature and status of the authorization dependency between the objects. For information, see *Classification of Authorization Dependencies Between Objects*.

Full information about the connection is also displayed in the *Properties* view when you select the connection.

i Note

If the *Properties* view is not visible, from the main menu choose ► *Window* ► *Show View* ► *Properties* ►.

2. Isolate the object(s) with missing authorization by choosing the  *Show missing authorization only* button.
3. Optional: If necessary, manipulate the view to help your analysis using the available toolbar options.
4. Grant the missing privilege(s) to the user with the invalid dependency.

This might be your user if you are the object owner, but it might also be the owner of another object if you are facing a complex object hierarchy.
5. In the authorization dependency viewer, refresh () the view to verify the validity of previously invalid dependencies.

Related Information

[Classification of Authorization Dependencies Between Objects \[page 773\]](#)

7.2.4.10.1 Example: Resolving an Invalidated Procedure Error

This example shows you how you identify the source of an invalidated procedure error using the authorization dependency viewer.

Context

Assume the following:

User DEPVIEWER is the owner of the schema DEPVIEWER, which contains the objects DEPVIEW and DEPTABLE.

User BODOS creates the procedures PROC_TO_PROC_HIER, PROC_TO_PROC, and PROC_TO_DEPVIEWER. The objects are dependent on each other as follows:

- PROC_TO_PROC_HIER executes the procedures PROC_TO_DEPVIEWER and PROC_TO_PROC.
- PROC_TO_PROC executes PROC_TO_DEPVIEWER
- PROC_TO_PROC selects and deletes from DEPTVIEW.
- PROC_TO_DEPVIEWER selects from DEPTABLE and DEPTVIEW.
- DEPTVIEW selects from DEPTABLE.

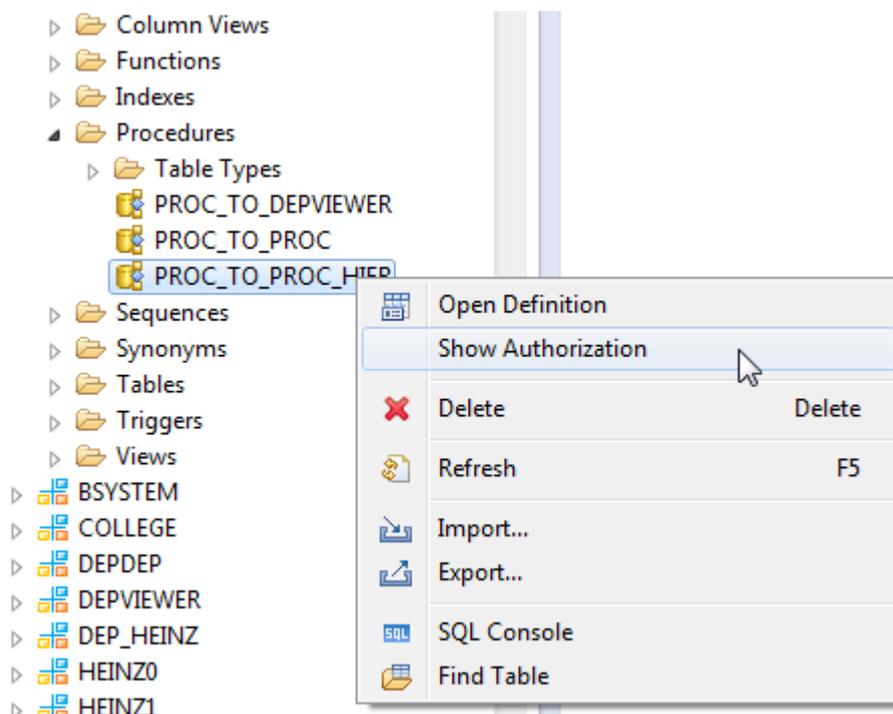
Other users are now granted EXECUTE privilege on PROC_TO_PROC_HIER. However, when they execute the procedure, the following error appears:

```
Could not execute 'call PROC_TO_PROC_HIER' SAP DBTech JDBC: [430]: invalidated procedure: PROC_TO_PROC_HIER: line 1 col 6 (at pos 5)
```

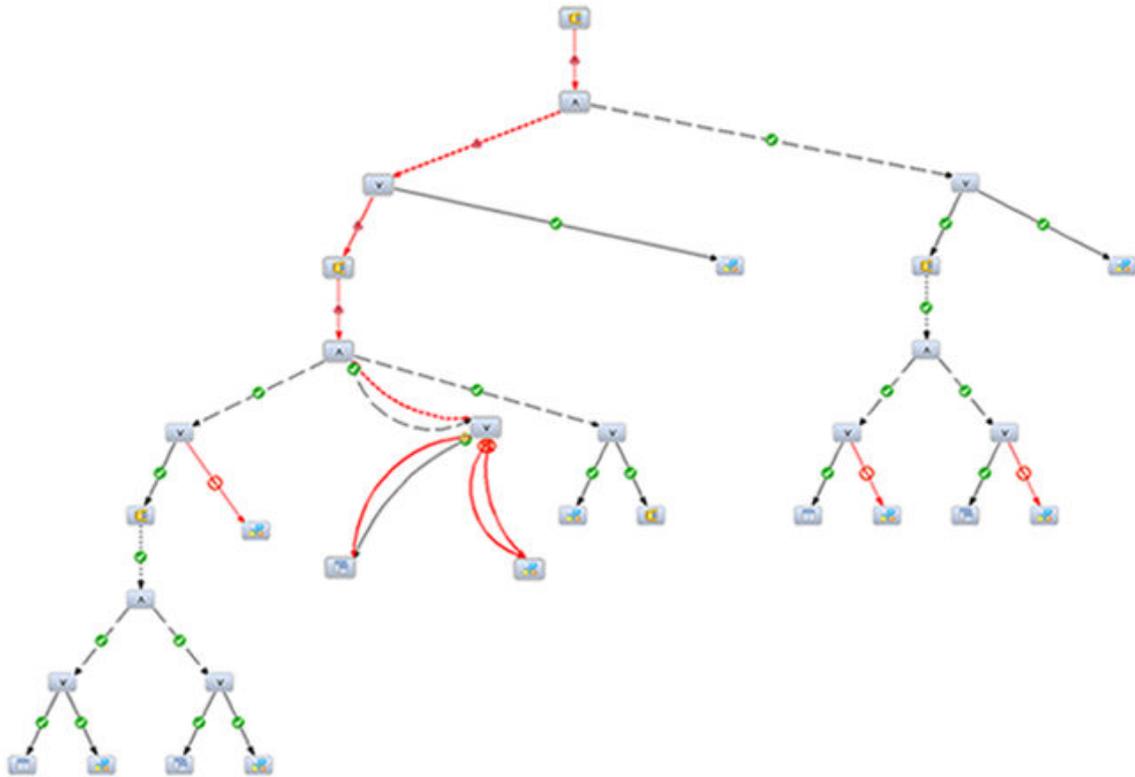
You can use the authorization dependency viewer to isolate the source of the problem as follows:

Procedure

1. In the *Systems* view, navigate to the procedure PROC_TO_PROC_HIER and from the context menu, choose *Show Authorization*:

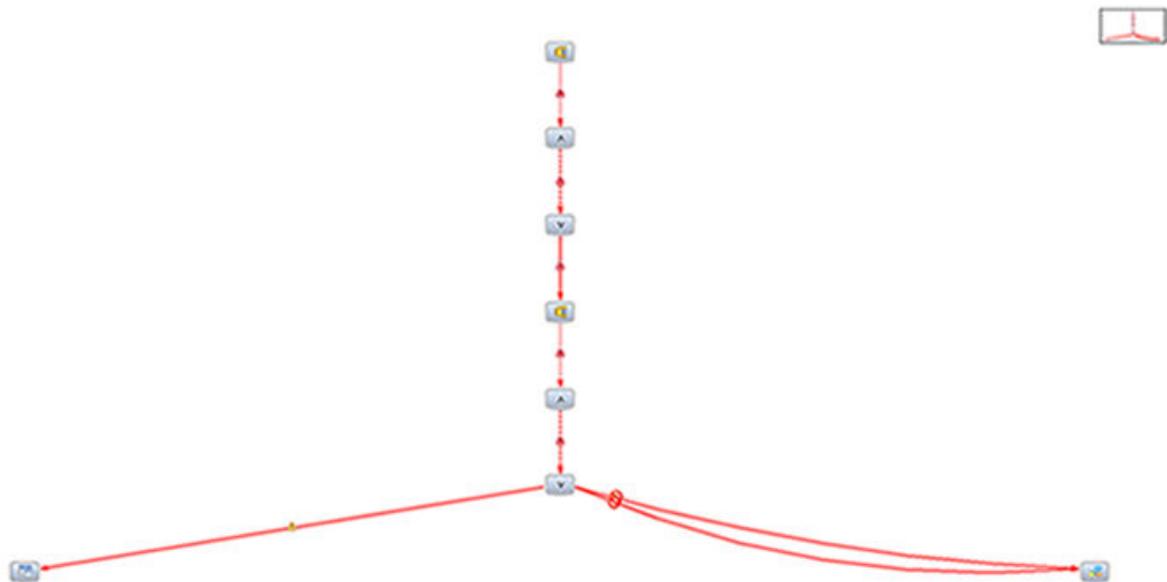


The full authorization dependency structure of the procedure is displayed as a hierarchical tree:

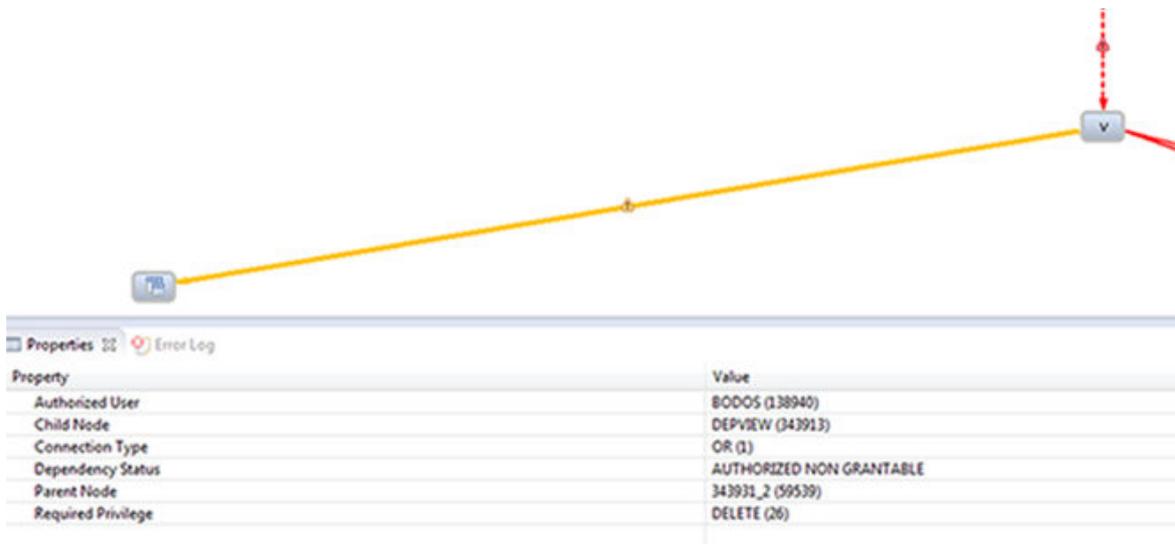


- From the toolbar, choose  (*Show missing authorization only*).

Only the invalid dependency path is shown. You can see that privileges are missing on either the view DEPVIEW or its parent schema DEPVIERER:

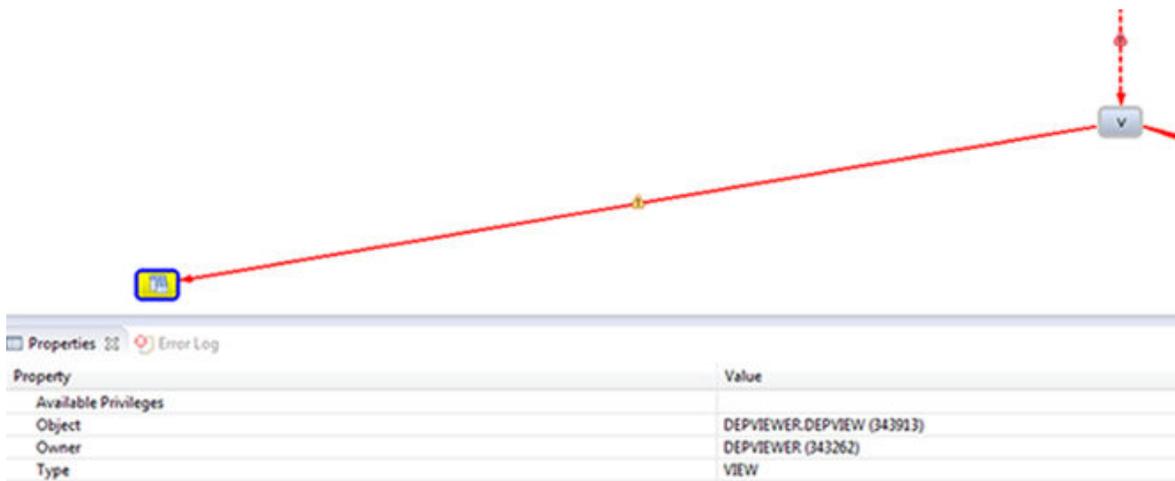


- To examine the invalid dependency path in more detail, select the connection to the view.



In the *Properties* view, you can see that the owner of the procedure has the required DELETE privilege on the underlying view, but is not authorized to grant this privilege further (dependency status is AUTHORIZED NON GRANTABLE). This invalidates the procedure that references the view.

- To see who owns the view (and therefore who needs to grant the missing authorization) select the object.



In the *Properties* view, you can see that the view DEPVIEW is owned by the user DEPVIEWER.

- As user DEPVIEWER, in the user definition of user BODOS, select *Grantable to others* for the EXECUTE privilege on the object DEPVIEW:

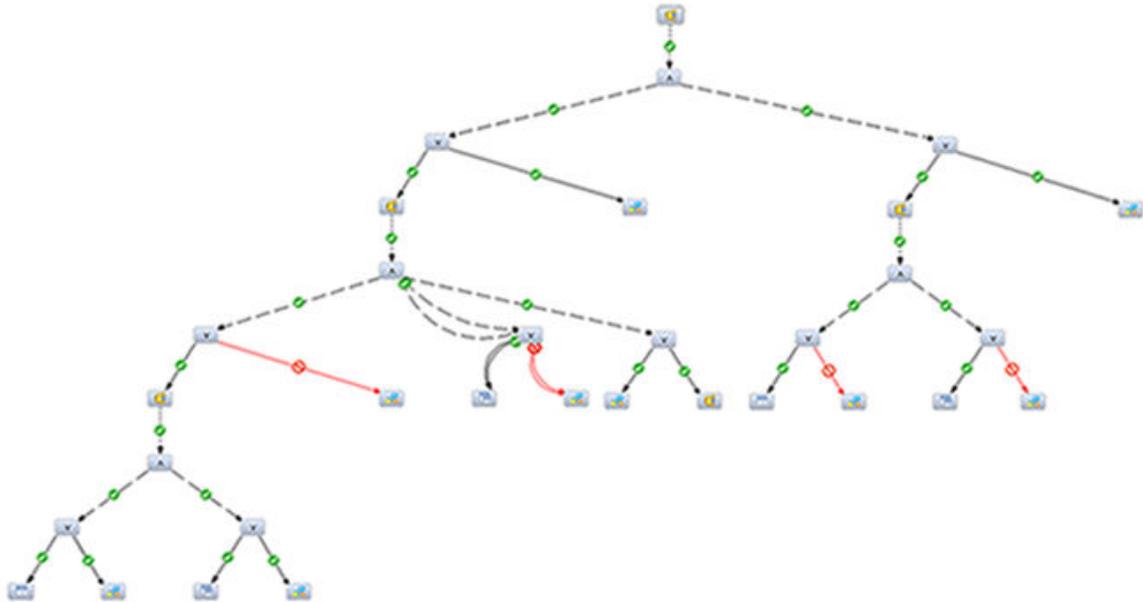


Note

Any user to whom user DEPVIEWER has granted the required privilege with authorization to grant further could also grant the missing authorization to user BODOS.

6. In the authorization dependency viewer, choose .

There are now no invalid authorization dependencies; the procedure is valid ():



7.2.4.10.2 Classification of Authorization Dependencies Between Objects

The authorization dependency viewer visualizes a root object's authorization dependency structure as a hierarchical tree. The lines connecting the nodes in the tree indicate the nature and status of the authorization dependency between the objects.

Connection	Description
Long dash line (----)	An AND connection exists between the parent node and the child nodes. Access to the parent node requires authorization to all child nodes.
Solid line (-----)	An OR connection exists between the parent node and the child nodes. Access to the parent node requires authorization to one of the child nodes.
Black line	The authorization dependency status is valid, that is, the user has the required privilege to the child object and is authorized to grant it further. This is additionally indicated by the  (AUTHORIZED GRANTABLE) icon.

Connection	Description
Red line	<p>The authorization dependency status is invalid in some way. The following icons indicate the exact status:</p> <ul style="list-style-type: none">  (NOT AUTHORIZED) The user does not have the required privilege for the child object.  (AUTHORIZED NON GRANTABLE) The user has the required privilege for the child object but is not authorized to grant it further because he is missing WITH GRANT OPTION.  (AUTHORIZED NON GRANTABLE_ENFORCED) The user has the required privilege for the child object but is not able to grant it further because it itself is not grantable. This fact determines the dependency status of the parent object even if the parent object has an OR connection to another child object with valid authorization.  (INVALID) The user does not have the required privilege for the child object or the child object is invalidated. This fact determines the dependency status of the parent object even if the parent object has an OR connection to another child object with a valid dependency status.

7.2.4.10.3 Toolbar Options in the Authorization Dependency Viewer

Several options in the authorization dependency viewer allow you to manipulate the view to help your analysis of authorization errors.

Option	Description
 (Switch to the graph view)	<p>Opens the graph view</p> <p>This view shows the dependency structure as a graph. In the tree view, the same database object might appear multiple times if it is referenced at different levels of the tree. In the graph view, each database object is only one node. This feature might be helpful in identifying the single root cause of your problem.</p>
 (Switch to the object dependencies only view)	<p>Opens the object dependencies view</p> <p>This view shows the transitive closure of all objects on which the view or procedure depends. This tree does not contain duplicate nodes or meta nodes.</p>
 (Zoom in) /  (Zoom out)	<p>Zooms in or out of the dependency structure for the required level of detail</p>

Option	Description
 (Reset zoom)	Resets the view after zooming
 (Auto arrange)	Resets the view after rearranging

7.2.5 Provisioning Users

As a user administrator, you create and configure database users, as well as authorize them to work with the SAP HANA database.

The recommended process for provisioning users is as follows:

1. Define and create roles.
2. Create users.
3. Grant roles to users.

Further tasks related to user provisioning include for example:

- Deleting users when they leave the organization
- Reactivating users after too many failed logon attempts
- Deactivating users if a security violation has been detected
- Resetting user passwords

i Note

If you are using an Identity Management (IDM) system for user provisioning, it is highly recommended that you create a dedicated technical user for that system that has the system privileges USER ADMIN and ROLE ADMIN and object privilege EXECUTE on the procedure GRANT_ACTIVATED_ROLE. This database user should then be used exclusively by the IDM system for its user provisioning tasks.

[Provisioning Users in SAP HANA Cockpit \[page 776\]](#)

You can use the SAP HANA cockpit to create database users and grant them roles.

[Provisioning Users in SAP HANA Studio \[page 806\]](#)

You can use the *User* and *Role* editors of the SAP HANA studio to perform user-provisioning tasks.

[Provisioning Users Using an LDAP Identity Management Server \[page 818\]](#)

The Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing directory services. If you use an LDAP-compliant directory server to manage users and their access to resources, you can leverage LDAP group membership to authorize users.

Related Information

[Catalog Roles and Design-Time Roles Compared \[page 761\]](#)

7.2.5.1 Provisioning Users in SAP HANA Cockpit

You can use the SAP HANA cockpit to create database users and grant them roles.

i Note

It's not possible by using the cockpit to create the technical user required to register a resource in the SAP HANA cockpit. You need to create this user and grant the minimum necessary authorization by using SQL as follows:

```
CREATE USER <username> PASSWORD <password> NO FORCE_FIRST_PASSWORD_CHANGE;  
GRANT CATALOG READ to <username>;  
GRANT SELECT on SCHEMA _SYS_STATISTICS to <username>
```

[View a Database User \[page 777\]](#)

You can view database users on the *User* page of the SAP HANA cockpit.

[View a Database Role \[page 783\]](#)

You can view database roles on the *Role* page of the SAP HANA cockpit.

[Create a Catalog Role \[page 786\]](#)

You can create a new role directly in runtime and grant it the privileges and roles necessary for the task or function that it represents on the *Role* page of the SAP HANA cockpit. It is also possible to map roles to LDAP groups if you are implementing user authorization based on LDAP group membership.

[Change a Role \[page 790\]](#)

You can change the roles and privileges assigned to a role on the *Role* page of the SAP HANA cockpit.

[Delete a Role \[page 791\]](#)

You can delete a role on the *Role* page of the SAP HANA cockpit.

[Create a Database User \[page 792\]](#)

You create a standard database user for every person who needs to work directly with the SAP HANA database. When you create a user, you also configure how the user will be authenticated. You can do this on the *User* page of the SAP HANA cockpit.

[Create a Restricted Database User \[page 796\]](#)

You create a restricted user for users who access SAP HANA through client applications – full SQL access via an SQL console is not intended. When you create a restricted user, you also configure how the user will be authenticated. You can do this on the *User* page of the SAP HANA cockpit.

[Assign Roles to a Database User \[page 799\]](#)

Roles are the standard mechanism of granting privileges to SAP HANA database users. It is recommended that you assign roles to users instead of granting privileges individually. You can grant roles to users on the *Assign Roles* page of the SAP HANA cockpit.

[Assign Privileges to a User \[page 801\]](#)

It is recommended that you assign roles to users instead of granting privileges individually. However, you can still grant privileges directly to users using the *Assign Privileges* app.

[Change a Database User \[page 802\]](#)

You can change an existing database user on the *User* page of the SAP HANA cockpit.

[Deactivate a Database User \[page 803\]](#)

Users can be automatically deactivated for security reasons, for example, if they violate password policy rules. However, as a user administrator, you may need to explicitly deactivate a user, for example,

if an employee temporarily leaves the company or a security violation is detected. You can deactivate a user on the [User](#) page of the SAP HANA cockpit.

[Delete a Database User \[page 804\]](#)

You may need to delete a database user if an employee leaves your organization for example. You can delete a user with on the [User](#) page of the SAP HANA cockpit.

[Add a SAML Identity Provider in SAP HANA Cockpit \[page 733\]](#)

If you are implementing Security Assertion Markup Language (SAML) to authenticate users accessing SAP HANA via the SQL interface directly (that is using JDBC and ODBC clients), you must add the SAML identity providers for the required users. You can do this using the SAP HANA cockpit.

7.2.5.1.1 View a Database User

You can view database users on the [User](#) page of the SAP HANA cockpit.

Prerequisites

You have the system privilege CATALOG READ. You don't require any additional privileges to view your own database user.

Procedure

1. On the [Overview](#) page, choose the [Manage users](#) link.
The [User](#) page opens. All existing database users are displayed in list format on the left.
2. To see more detailed information about a specific user, simply select it.
For more information, see [Database User Details](#).

7.2.5.1.1.1 Database User Details

On the [User](#) page of the SAP HANA cockpit, you can view the details of all users in the SAP HANA database.

General Information

Field	Description
User Name	Unique user name

Field	Description
<i>E-Mail</i>	User's e-mail address
<i>Valid From/To</i>	<p>Validity period of the user</p> <p>If the user account is not currently within its validity period, the user is inactive and cannot log on.</p> <p>If no validity period is configured, the user is indefinitely valid.</p>
<i>Creation of Objects in Own Schema</i>	<p>Indicates whether or not the user can create objects in their database schema</p> <p>Standard users can create objects in their schema. Restricted users cannot.</p> <p>For more information about the difference between standard and restricted users, see <i>Database Users</i>.</p>
<i>PUBLIC Role</i>	<p>Indicates whether or not the user has the PUBLIC role</p> <p>Standard users have this role by default. Restricted users do not.</p> <p>For more information about the difference between standard and restricted users, see <i>Database Users</i>.</p>
<i>Disable ODBC/JDBC Access</i>	<p>Indicates whether or not the user can connect to the database via ODBC or JDBC</p> <p>By default, ODBC/JDBC access is disabled for restricted users, meaning they can only connect via HTTP/HTTPS, and enabled for standard users.</p> <p>For more information about the difference between standard and restricted users, see <i>Database Users</i>.</p>
<i>Comment</i>	Free-text comment or description (if applicable)

Authorization Mode

Field	Description
Authorization Mode	<p>Indicates whether the user's authorization is based on LDAP group membership or local SAP HANA mechanisms</p> <p>A user with authorization mode LDAP is granted roles exclusively based on their LDAP group membership. It is not possible to grant such a user other roles or privileges directly.</p> <p>The default user authorization mode is Local. This means that the user must be granted roles and privileges directly.</p> <p>For more information about LDAP group authorization, see the <i>SAP HANA Security Guide</i>.</p>
Assign Roles	Open the Assign Roles app where you can assign roles to the user
Assign Privileges	Open the Assign Privileges app where you can assign privileges to the user

Authentication

Field	Description
Password	<p>Indicates whether or not user name password authentication is enabled</p> <p>Users accessing the SAP HANA database authenticate themselves by entering their database user name and their local SAP HANA password.</p>
Force password change on next logon	Indicates whether the user must change a password set by a user administrator the first time he or she logs on, regardless of how the password policy parameter Password Change Required on First Logon is configured

Field	Description
<i>Kerberos</i>	<p>Indicates whether or not Kerberos authentication is enabled</p> <p>If enabled, the external identity to which the database user is mapped must be specified.</p> <p>A Kerberos authentication provider can be used to authenticate users accessing SAP HANA in the following ways:</p> <ul style="list-style-type: none"> • Directly from ODBC and JDBC database clients within a network (for example, the SAP HANA studio) • Indirectly from front-end applications such as SAP BusinessObjects applications and other SAP HANA databases using Kerberos delegation • Via HTTP/HTTPS access by means of SAP HANA Extended Services (SAP HANA XS), advanced model and classic model <p>In this case, Kerberos authentication is enabled with Simple and Protected GSSAPI Negotiation Mechanism (SPNEGO).</p>
<i>SAP Logon Ticket, SAP Assertion Ticket</i>	<p>Indicates whether or not authentication using SAP logon or assertion tickets is enabled</p> <p>Users can be authenticated by SAP logon or assertion tickets issued to them when they log on to an SAP system that is configured to create tickets (for example, the SAP Web Application Server or Portal).</p> <div data-bbox="804 1160 1394 1330" style="background-color: #f0f0f0; padding: 10px;"> <p>i Note</p> <p>To implement logon/assertion tickets, the user specified in the logon/assertion ticket must already exist in SAP HANA; there is no support for user mapping.</p> </div>
<i>SAML</i>	<p>Indicates whether or not SAML (security assertion markup language) authentication is enabled</p> <p>If enabled, the external identity to which the database user is mapped can be explicitly specified. Alternatively, if the option <i>Automatic Mapping by Provider</i> is selected, the identity provider is allowed to map its users to the database user.</p> <p>A SAML bearer assertion can be used to authenticate users accessing SAP HANA directly from ODBC/JDBC database clients. SAP HANA can act as a service provider to authenticate users accessing via HTTP/HTTPS by means of SAP HANA XS classic and advanced.</p>

Field	Description
X509	<p>Indicates whether or not X.509 certificate authentication is enabled</p> <p>For HTTP/HTTPS access to SAP HANA by means of SAP HANA XS advanced model and classic model, users can be authenticated by client certificates signed by a trusted Certification Authority (CA), which can be stored in the SAP HANA XS trust store.</p> <div style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>To implement X.509 client certificates, the user specified in the certificate must already exist in SAP HANA; there is no support for user mapping.</p> </div>
JWT	<p>Indicates whether or not JSON Web Token authentication is enabled</p> <p>If enabled, the external identity to which the database user is mapped can be explicitly specified. Alternatively, if the option <i>Automatic Mapping by Provider</i> is selected, the identity provider is allowed to map its users to the database user.</p> <p>A JSON Web Token can be used to authenticate users accessing SAP HANA directly from ODBC/JDBC database clients or indirectly through SAP HANA extended application services, advanced model (SAP HANA XS, advanced).</p>

Custom User Properties

Additional user properties can be configured for client applications. The following properties are available by default:

Property	Description
CLIENT	<p>The session client</p> <p>When you create SAP HANA calculation views, it is possible to filter the data according to the client specified in table fields such as MANDT or CLIENT.</p>
LOCALE	<p>The user's locale</p> <p>When you create SAP HANA information models (attribute views, analytic views, and calculation views), this parameter can be used to translate information according to the user's locale.</p>

Property	Description
PRIORITY	<p>The priority with which the thread scheduler handles statements executed by the user</p> <p>Priority values of 0 (lowest priority) to 9 (highest) are available; the default priority is 5.</p>
STATEMENT MEMORY LIMIT	<p>The maximum memory (in GB) that can be used by a statement executed by the user</p> <p>The properties <code>statement_memory_limit</code> and <code>statement_memory_limit_threshold</code> in the <code>memory_manager</code> section of the <code>global.ini</code> configuration file are used to limit the memory that can be allocated with respect to statement execution.</p> <p><code>statement_memory_limit_threshold</code> indicates what percentage of the global memory allocation limit must be in use before the specific value of <code>statement_memory_limit</code> is applied. If this memory limit is being applied and a statement execution exceeds it, then the statement is aborted.</p> <p>With this user parameter, you can set a user-specific limit that takes precedence over the global statement memory limit.</p> <p>For more information about memory usage, see <i>Setting a Memory Limit for SQL Statements</i> in the <i>SAP HANA Administration Guide</i>.</p>
STATEMENT THREAD LIMIT	<p>The maximum number of threads that can be used by a statement executed by the user</p>
TIME ZONE	<p>The user's timezone</p> <p>The standard database formats for locale and timezone are supported.</p>

Related Information

[Setting a Memory Limit for SQL Statements \[page 633\]](#)

[Configure LDAP Group Authorization \[page 819\]](#)

[Assign Roles to a Database User \[page 799\]](#)

[Assign Privileges to a User \[page 801\]](#)

7.2.5.1.2 View a Database Role

You can view database roles on the [Role](#) page of the SAP HANA cockpit.

Prerequisites

You have the system privilege CATALOG READ or ROLE ADMIN.

i Note

Even if ROLE ADMIN was revoked from your user, you can still view roles that you created yourself.

Procedure

1. On the [Overview](#) page, choose the [Manage roles](#) link.
The [Role](#) page opens. All existing database roles are displayed in list format on the left.
2. To see more detailed information about a specific role, simply select it.
For more information, see [Database Role Details](#).

7.2.5.1.2.1 Database Role Details

On the [Role](#) page of the SAP HANA cockpit, you can view the details of all roles in the SAP HANA database.

General Information

Field	Description
Schema	<p>Schema in which the role exists (if applicable)</p> <p>The schema represents the role's runtime namespace, which allows it to be used in different contexts.</p> <p>A role without a schema is a global role.</p>

⚠ Caution

A role with a namespace will be deleted if the schema is deleted.

Field	Description
<i>Creator</i>	User who created the role
<i>LDAP Groups</i>	<p>LDAP groups that have been mapped to the role (if applicable)</p> <p>Database users configured for LDAP authorization who belong to the specified group(s) are automatically granted the role in line with your LDAP configuration for SAP HANA. For more information, see the <i>SAP HANA Security Guide</i>.</p>
<i>Comment</i>	Free-text comment or description (if applicable)
<i>Type</i>	<p>The role type:</p> <ul style="list-style-type: none"> • <i>Catalog</i> A role created in run-time with the SQL statement <code>CREATE ROLE</code> • <i>Catalog (LDAP)</i> A catalog role with LDAP group mappings • <i>HDI</i> A role created using the SAP Web IDE for SAP HANA and deployed using SAP HANA deployment infrastructure (SAP HANA DI) • <i>HDI (LDAP)</i> A HDI role with LDAP group mappings LDAP groups are mapped to the activated catalog role. • <i>Repository role</i> A role created in the built-in repository of the SAP HANA database using either the SAP HANA Web Workbench or the SAP HANA studio
<i>Is Part of Roles</i>	Indicates whether the role is included in another role
<i>Roles</i>	Other roles granted to this role
<i>System Privileges</i>	<p>System privileges granted to the role</p> <p>System privileges control general system activities. They are mainly used for administrative purposes, such as creating schemas, creating and changing users and roles, performing data backups, managing licenses, and so on.</p> <p>For a list of all system privileges, see <i>System Privileges (Reference)</i>.</p>
<i>Object Privileges</i>	<p>Object privileges granted to the role</p> <p>Object privileges are used to allow access to and modification of database objects, such as tables and views. Depending on the object type, different actions can be authorized (for example, SELECT, CREATE ANY, ALTER, DROP, and so on).</p> <p>For a list of all object privileges, see <i>Object Privileges (Reference)</i>.</p>

Field	Description
<i>Analytic Privileges</i>	<p>Analytic privileges granted to the role</p> <p>Analytic privileges are used to control read access to data in SAP HANA information models (that is, analytic views, attribute views, and calculation views) depending on certain values or combinations of values.</p>
<i>Is Part of Roles</i>	Other roles to which this role has been granted
<i>Application Privileges</i>	<p>Application privileges granted to the role</p> <p>Application privileges are used to authorize user and client access to SAP HANA XS classic applications.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>Application privileges are not relevant in the context of SAP HANA XS advanced applications. For more information about the authorization concept of the SAP HANA XS advanced, see the <i>SAP HANA Security Guide</i>.</p> </div>
<i>Package Privileges</i>	<p>Package privileges granted to the role</p> <p>Package privileges are used to allow access to and the ability to work in packages in the repository of the SAP HANA database.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>With SAP HANA XS advanced, source code and web content are not versioned and stored in the repository, so package privileges are not relevant in this context.</p> </div>
<i>Privileges on Users</i>	<p>Privileges on users granted to the role</p> <p>ATTACH DEBUGGER is the only privilege that can be granted on a user.</p> <p>For example, User A can grant User B the privilege ATTACH DEBUGGER to allow User B debug SQLScript code in User A's session. User A is only user who can grant this privilege. Note that User B also needs the object privilege DEBUG on the relevant SQLScript procedure.</p>

Related Information

[Database Roles \[page 759\]](#)

[System Privileges \[page 742\]](#)

[System Privileges \(Reference\) \[page 743\]](#)

[Object Privileges \[page 748\]](#)

[Object Privileges \(Reference\) \[page 749\]](#)

[Analytic Privileges \[page 754\]](#)

[Package Privileges \[page 756\]](#)

[Application Privileges \[page 758\]](#)

7.2.5.1.3 Create a Catalog Role

You can create a new role directly in runtime and grant it the privileges and roles necessary for the task or function that it represents on the [Role](#) page of the SAP HANA cockpit. It is also possible to map roles to LDAP groups if you are implementing user authorization based on LDAP group membership.

Prerequisites

- You have the system privilege `ROLE ADMIN`.
- You have the privileges required to grant privileges and roles to the new role. For more information, see *Prerequisites for Granting and Revoking Privileges and Roles*.

Procedure

1. On the [Overview](#) page, choose the [Manage roles](#) link.

The [Role](#) page opens. All existing database roles are displayed in list format on the left.

2. Create the role:
 - a. Click the **+Add** button in the footer toolbar.
 - b. Specify a unique role name.
 - c. Optional: Enter a comment or text to describe the role.
 - d. Optional: Assign the role a runtime namespace by choosing the schema in which to create the role.

Role namespaces allow you to reuse role names in different contexts. If you do not select a schema, the role will be created as a global role.

Caution

A role with a namespace will be deleted if the schema is deleted.

3. Optional: If you are implementing user authorization based on LDAP group membership, map one or more LDAP group to the role:
 - a. Enable the assignment of LDAP groups to the role.
 - b. Add the required LDAP groups by specifying the unique distinguished name (DN).

Users configured for LDAP authorization who belong to the specified group(s) are automatically granted the role in line with your LDAP configuration for SAP HANA. For more information, see the *SAP HANA Security Guide*.

4. Choose [Save](#).

The role is created.

5. Assign the required roles to the role:
 - a. In the *Roles* area, choose *Edit*.
 - b. Choose *Add* and select the roles you want to assign.
 - c. If you want users who have the new role to be able to grant the assigned role on to others, choose *Grantable to Others*.
 - d. Save the role assignment.
6. Assign the required privileges to the role:
 - a. For the relevant privilege type, choose *Edit*.
 - b. Choose *Add* and select the privileges you want to assign.

i Note

For object and package privileges, you must first add the object or package and then add the required privilege to the object or package.

- c. If you want users who have the new role to be able to grant the assigned privilege on to others, choose *Grantable to Others*.
- d. Save the privilege assignment.
- e. Repeat for further privilege types.

Results

The role is created and appears in the list of roles on the left.

Next Steps

Assign the role to the required database users. You should only do this for users with authorization mode *Local*, not *LDAP*.

If you mapped LDAP groups to the role, configure the connection to the LDAP provider and configure the required database users for LDAP group authorization. For more information, see the *SAP HANA Administration Guide*.

Related Information

[Privileges \[page 740\]](#)

[System Privileges \(Reference\) \[page 743\]](#)

[Object Privileges \(Reference\) \[page 749\]](#)

[Database Role Details \[page 783\]](#)

[Assign Roles to a Database User \[page 799\]](#)

[Configure LDAP Group Authorization \[page 819\]](#)

7.2.5.1.3.1 Prerequisites for Granting and Revoking Privileges and Roles

To be able to grant and revoke privileges and roles to and from users and roles, several prerequisites must be met.

The following table lists the prerequisites that a user must meet to grant privileges and roles to another user (or role).

Prerequisites for Granting Privileges

To grant...	The granting user needs...
A system privilege	The system/object privilege being granted and be authorized to grant it to other users and roles
An object privilege on an object that exists only in runtime	
An object privilege on an activated object created in the repository, such as a calculation view	The object privilege EXECUTE on the procedure GRANT_PRIVILEGE_ON_ACTIVATED_CONTENT
An object privilege on schema containing activated objects created in the repository, such as a calculation view	The object privilege EXECUTE on the procedure GRANT_SCHEMA_PRIVILEGE_ON_ACTIVATED_CONTENT
A package privilege	The package privilege being granted and be authorized to grant it to other users and roles
An analytic privilege	The object privilege EXECUTE on the procedure GRANT_ACTIVATED_ANALYTICAL_PRIVILEGE
An application privilege	The object privilege EXECUTE on the procedure GRANT_APPLICATION_PRIVILEGE
Privilege on user ATTACH DEBUGGER	To be the user on which ATTACH DEBUGGER is granted
A role created in runtime	Either: <ul style="list-style-type: none"> • The role being granted and be authorized to grant it to other users and roles, or • The system privilege ROLE ADMIN
A role created in the repository	The object privilege EXECUTE on the procedure GRANT_ACTIVATED_ROLE
A role created in an HDI container	Privileges to execute GRANT_CONTAINER_SCHEMA_ROLES in the container's API schema, or, if the user is a container group administrator, privileges to execute GRANT_CONTAINER_SCHEMA_ROLES in the container group's API schema

Prerequisites for Revoking Privileges

To revoke ...	The revoking user needs...
A system privilege	To be the user who granted the privilege
An object privilege on an object that exists only in runtime	
An object privilege on an activated object created in the repository, such as a calculation view	The object privilege EXECUTE on the procedure REVOKE_PRIVILEGE_ON_ACTIVATED_CONTENT

To revoke ...	The revoking user needs...
An object privilege on schema containing activated objects created in the repository, such as a calculation view	The object privilege EXECUTE on the procedure REVOKE_SCHEMA_PRIVILEGE_ON_ACTIVATED_CONTENT
A package privilege	The user who granted the privilege
An analytic privilege	The object privilege EXECUTE on the procedure REVOKE_ACTIVATED_ANALYTICAL_PRIVILEGE
An application privilege	The object privilege EXECUTE on the procedure REVOKE_APPLICATION_PRIVILEGE
Privilege on user ATTACH DEBUGGER	To be the user on which ATTACH DEBUGGER is granted
A role created in runtime	<ul style="list-style-type: none"> To be the user who granted the role, or The system privilege ROLE ADMIN
<div style="border: 1px solid #ccc; padding: 5px; background-color: #f0f0f0;"> <p>i Note</p> <p>With the exception of roles granted by technical user <code>_SYS_REPO</code>, a user with <code>ROLE ADMIN</code> cannot revoke roles granted by technical users <code>SYS</code> and <code>_SYS*</code>.</p> </div>	
A role created in the repository	The object privilege EXECUTE on the procedure REVOKE_ACTIVATED_ROLE
A role created in an HDI container	Privileges to execute <code>REVOKE_CONTAINER_SCHEMA_ROLES</code> in the container's API schema, or, if the user is a container group administrator, privileges to execute <code>REVOKE_CONTAINER_SCHEMA_ROLES</code> in the container group's API schema

Authorization of User `_SYS_REPO`

If you are implementing your authorization concept using roles and you are creating roles in the repository of the SAP HANA database, the technical user `_SYS_REPO` is the granting and revoking user. `_SYS_REPO` **automatically** meets all of the above prerequisites with the exception of those for granting/revoking objects privileges on objects that exist only in runtime. These privileges must be explicitly granted to `_SYS_REPO`. For more information about roles as repository objects, see the *SAP HANA Security Guide*.

How are HDI roles granted and revoked?

In the SAP HANA deployment infrastructure (HDI), there are two types of users that can grant or revoke roles.

Roles created in an HDI container can be granted to (or revoked from) other roles and users either by an administrator of the container or an administrator of the container group that the container belongs to.

Preferably, roles created in an HDI container are granted or revoked by a container administrator.

Related Information

[Create and Authorize a User \[page 808\]](#)

[Create and Authorize a Restricted User \[page 810\]](#)

7.2.5.1.4 Change a Role

You can change the roles and privileges assigned to a role on the [Role](#) page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege ROLE ADMIN or are the owner of the role.
- You have the privileges required to grant privileges and roles to the role. For more information, see *Prerequisites for Granting and Revoking Privileges and Roles*.

Context

It is possible to change a catalog role by revoking roles and privileges from the role or granting further roles and privileges.

⚠ Caution

Do not change roles that were originally created in design time, that is, **HDI roles** or **repository roles**. If you change the runtime version of such a role, your changes will be overwritten the next time a new version of the design-time role is deployed. For more information about creating roles in design time, see the SAP HANA developer documentation.

Procedure

1. On the [Overview](#) page, choose the [Manage roles](#) link.

The [Role](#) page opens. All existing database roles are displayed in list format on the left.

2. Find the role you want to change.

→ Tip

Search for a role by entering the name or part of the name in the search box.

3. To change the role's LDAP group mapping or comment, choose the [Edit](#) button in the header area.
4. To change the roles or privileges assigned to the role, select the relevant tab page and choose [Edit](#).

5. Make the required changes and save.

Related Information

[Database Role Details \[page 783\]](#)

7.2.5.1.5 Delete a Role

You can delete a role on the [Role](#) page of the SAP HANA cockpit.

Prerequisites

You have the system privilege ROLE ADMIN.

Procedure

1. On the [Overview](#) page, choose the [Manage roles](#) link.
The [Role](#) page opens. All existing database roles are displayed in list format on the left.
2. Find the role you want to change.

→ Tip

Search for a role by entering the name or part of the name in the search box.

3. Choose [Delete](#).

Results

The role is deleted.

i Note

You can also use the above procedure to delete HDI and repository roles. However, these roles will be recreated when they are deployed again.

7.2.5.1.6 Create a Database User

You create a standard database user for every person who needs to work directly with the SAP HANA database. When you create a user, you also configure how the user will be authenticated. You can do this on the [User](#) page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege `USER ADMIN`.
- If you are integrating SAP HANA database users into a single sign-on (SSO) environment using one or more of the supported mechanisms, the necessary infrastructure must be in place and configured. For more information about SSO, see the *SAP HANA Security Guide*.
- If you are implementing LDAP group authorization, the necessary infrastructure must be in place and configured. For more information, see the section on configuring LDAP group authorization in the *SAP HANA Administration Guide*.

Procedure

1. On the [Overview](#) page, choose the [Manage users](#) link.

The [User](#) page opens. All existing database users are displayed in list format on the left.

2. Create a new user by clicking the **+Add** button in the footer toolbar and choosing [Create User](#).
3. Specify the new user name.

You must give the user a unique name. User names can contain any CESU-8 characters except for a small subset. For more information, see *Unpermitted Characters in User Names*.

4. Optional: Specify the user's e-mail address.
5. Optional: Specify a validity period for the user, including the appropriate time zone.

For example, if you are creating a user for a new employee, you can enter their start date in the [Valid From](#) field.

If you do not enter any values, the user is immediately and indefinitely valid.

6. Optional: Prevent the user from being able to create objects in his own database schema by selecting [No](#) for the option [Creation of Objects in Own Schema](#).

i Note

If you select [No](#) for both this option **and** the next option ([PUBLIC Role](#)), the user will be created as a restricted user, not a standard user.

7. Optional: Prevent the user from being granted the standard PUBLIC role by selecting [No](#) for the option [PUBLIC Role](#).

The PUBLIC role contains the privileges for filtered read-only access to the system views. To see data in a particular view, the user also needs the SELECT privilege on the view.

i Note

If you select *No* for both this option **and** the previous option (*Creation of Objects in Own Schema*), the user will be created as a restricted user, not a standard user.

- Optional: Prevent the user from being able to connect to the database via ODBC and JDBC clients by selecting the corresponding checkbox.

By default, standard users have access via ODBC and JDBC clients.

If you disable ODBC/JDBC client access, the user can still connect via HTTP. Furthermore, disabling ODBC/JDBC access does not affect the user's authorizations or prevent the user from executing SQL commands via channels other than JDBC/ODBC.

- Optional: Enter a comment or text to describe the user.
- Optional: Set the authorization mode to LDAP if the user's authorization is based on LDAP group membership.

A user with authorization mode *LDAP* is granted roles exclusively based on their LDAP group membership. It is not possible to grant such a user other roles or privileges directly.

The default user authorization mode is *Local*. This means that the user must be granted roles and privileges directly as normal.

i Note

Setting the authorization mode of the user is only one step in the configuration of LDAP group authorization. For more information, see the section on configuring LDAP group authorization in the *SAP HANA Administration Guide*.

- Specify how the user can be authenticated.

i Note

You must specify at least one authentication mechanism. For more information about the supported mechanisms, see *Database User Details*.

Authentication Mechanism	Required Configuration
User name and password	Enter and confirm the user's initial password. You can override the password policy setting (<i>Password Change Required on First Logon</i>) that forces users to change a password set by a user administrator the first time they log on. This is useful for technical users, for example.
Kerberos	Enter the user principal name (UPN) specified in the Microsoft Active Directory or the Kerberos Key Distribution Center as the external ID.
SAP logon and assertion tickets	No additional user configuration required in user definition
SAML	Choose <i>Add Identity Provider</i> , select the identity provider, and then enter the user ID known to the SAML identity provider.

Authentication Mechanism	Required Configuration
	<p>Alternatively, you can allow the identify provider to map its users to the database user by enabling automatic mapping by provider.</p> <div data-bbox="826 456 1402 618" style="background-color: #f0f0f0; padding: 5px;"> <p>i Note</p> <p>The identity provider must already be created. You can do this on the SAML Identity Provider page or using the SQL statement <code>CREATE SAML PROVIDER</code>.</p> </div>
<p>JWT (JSON Web Token)</p>	<p>Choose Add Identity Provider, select the identity provider, and then enter the user ID known to the JWT identity provider.</p> <p>Alternatively, you can allow the identify provider to map its users to the database user by enabling automatic mapping by provider.</p> <div data-bbox="826 860 1402 1048" style="background-color: #f0f0f0; padding: 5px;"> <p>i Note</p> <p>The identity provider must already be created. You can do this using the SQL statement <code>CREATE JWT PROVIDER</code>. For more information, see the SAP HANA Administration Guide.</p> </div>
<p>X.509 certificate</p>	<p>Choose Add X509 Certificate Manually and enter the user's public key certificate information.</p> <div data-bbox="826 1144 1402 1272" style="background-color: #f0f0f0; padding: 5px;"> <p>i Note</p> <p>X.509 certificates are supported only for HTTP access through the SAP HANA XS classic server.</p> </div>

12. Optional: Specify additional user properties required by client applications.
You can select from the available properties (see *Database User Details*) or manually enter a property.
13. Save the user.

Results

The user is created and appears in the list of users on the left. A new schema is created for the user in the catalog. It has the same name as the user.

Next Steps

Assign roles or privileges to the user (authorization mode *Local* only).

Related Information

[Configure LDAP Group Authorization \[page 819\]](#)

[Unpermitted Characters in User Names \[page 795\]](#)

[Database User Details \[page 777\]](#)

[Assign Roles to a Database User \[page 799\]](#)

[Assign Privileges to a User \[page 801\]](#)

[Database Users \[page 707\]](#)

[Add a SAML Identity Provider in SAP HANA Cockpit \[page 733\]](#)

7.2.5.1.6.1 Unpermitted Characters in User Names

User names can contain any CESU-8 characters except for a small subset.

The following characters are not allowed as user names:

Unicode Character	Character	Name
U+0021	!	Exclamation mark
U+0022	"	Quotation mark
U+0024	\$	Dollar sign
U+0025	%	Percent sign
U+0027	'	Apostrophe
U+0028	(Left parenthesis
U+0029)	Right parenthesis
U+002A	*	Asterisk
U+002B	+	Plus sign
U+002C	,	Comma
U+002E	.	Full stop
U+002F	/	Solidus
U+003A	:	Colon
U+003B	;	Semicolon
U+003C	<	Less-than sign
U+003D	=	Equals sign
U+003E	>	Greater-than sign
U+003F	?	Question mark
U+0040	@	Commercial at
U+005B	[Left square bracket

Unicode Character	Character	Name
U+005C	\	Reverse solidus
U+005D]	Right square bracket
U+005E	^	Circumflex accent
U+0060	`	Grave accent
U+007B	{	Left curly bracket
U+007C		Vertical line
U+007D	}	Right curly bracket
U+007E	~	Tilde

7.2.5.1.7 Create a Restricted Database User

You create a restricted user for users who access SAP HANA through client applications – full SQL access via an SQL console is not intended. When you create a restricted user, you also configure how the user will be authenticated. You can do this on the [User](#) page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege `USER ADMIN`.
- If you are integrating SAP HANA database users into a single sign-on (SSO) environment using one or more of the supported mechanisms, the necessary infrastructure must be in place and configured. For more information about SSO, see the *SAP HANA Security Guide*.
- If you are implementing LDAP group authorization, the necessary infrastructure must be in place and configured. For more information, see the section on configuring LDAP group authorization in the *SAP HANA Administration Guide*.

Procedure

1. On the [Overview](#) page, choose the [Manage users](#) link.

The [User](#) page opens. All existing database users are displayed in list format on the left.

2. Create a new restricted user by clicking the **+Add** button in the footer toolbar and choosing [Create Restricted User](#).
3. Specify the new user name.

You must give the user a unique name. User names can contain any CESU-8 characters except for a small subset. For more information, see *Unpermitted Characters in User Names*.

4. Optional: Specify the user's e-mail address.

- Optional: Specify a validity period for the user, including the appropriate time zone.

For example, if you are creating a user for a new employee, you can enter their start date in the *Valid From* field.

If you do not enter any values, the user is immediately and indefinitely valid.

- Optional: Allow the user to create objects in her own database schema by selecting *Yes* for the option *Creation of Objects in Own Schema*.

i Note

If you select *Yes* for this option, the user will be created as a standard user, not a restricted user.

- Optional: Grant the user the standard PUBLIC role by selecting *Yes* for the corresponding option.

The PUBLIC role contains the privileges for filtered read-only access to the system views. To see data in a particular view, the user also needs the SELECT privilege on the view.

i Note

If you select *Yes* for this option, the user will be created as a standard user, not a restricted user.

- Optional: Allow the user to connect to the database via ODBC and JDBC clients by deselecting the corresponding checkbox.

By default, restricted users are only able to connect to the database using HTTP. You must explicitly allow access via ODBC and JDBC clients by changing this setting.

For full access to ODBC or JDBC functionality, you must grant restricted users the standard role `RESTRICTED_USER_ODBC_ACCESS` or `RESTRICTED_USER_JDBC_ACCESS`. You can do this on the *Assign Roles* page.

- Optional: Enter a comment or text to describe the user.
- Optional: Set the authorization mode to LDAP if the user's authorization is based on LDAP group membership.

A user with authorization mode *LDAP* is granted roles exclusively based on their LDAP group membership. It is not possible to grant such a user other roles or privileges directly.

The default user authorization mode is *Local*. This means that the user must be granted roles and privileges directly as normal.

i Note

Setting the authorization mode of the user is only one step in the configuration of LDAP group authorization. For more information, see the section on configuring LDAP group authorization in the *SAP HANA Administration Guide*.

- Specify how the user can be authenticated.

i Note

You must specify at least one authentication mechanism. For more information about the supported mechanisms, see *Database User Details*.

Authentication Mechanism	Required Configuration
User name and password	<p>Enter and confirm the user's initial password.</p> <p>You can override the password policy setting (Password Change Required on First Logon) that forces users to change a password set by a user administrator the first time they log on. This is useful for technical users, for example.</p>
Kerberos	<p>Enter the user principal name (UPN) specified in the Microsoft Active Directory or the Kerberos Key Distribution Center as the external ID.</p>
SAP logon and assertion tickets	<p>No additional user configuration required in user definition</p>
SAML	<p>Choose Add Identity Provider, select the identity provider, and then enter the user ID known to the SAML identity provider.</p> <p>Alternatively, you can allow the identify provider to map its users to the database user by enabling automatic mapping by provider.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>The identity provider must already be created. You can do this on the SAML Identity Provider page or using the SQL statement <code>CREATE SAML PROVIDER</code>.</p> </div>
JWT (JSON Web Token)	<p>Choose Add Identity Provider, select the identity provider, and then enter the user ID known to the JWT identity provider.</p> <p>Alternatively, you can allow the identify provider to map its users to the database user by enabling automatic mapping by provider.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>The identity provider must already be created. You can do this using the SQL statement <code>CREATE JWT PROVIDER</code>. For more information, see the <i>SAP HANA Administration Guide</i>.</p> </div>
X.509 certificate	<p>Choose Add X509 Certificate Manually and enter the user's public key certificate information.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>X.509 certificates are supported only for HTTP access through the SAP HANA XS classic server.</p> </div>

12. Optional: Specify additional user properties required by client applications.

You can select from the available properties (see *Database User Details*) or manually enter a property.

13. Save the user.

Results

The user is created and appears in the list of users on the left. A new schema is also created for the user in the catalog. It has the same name as the user. However, as a restricted user, the user is not authorized to create objects in this schema. For more information about all restrictions, see *Database Users*.

Next Steps

Assign roles to the user (authorization mode *Local* only).

Related Information

[Configure LDAP Group Authorization \[page 819\]](#)

[Unpermitted Characters in User Names \[page 795\]](#)

[Database User Details \[page 777\]](#)

[Assign Roles to a Database User \[page 799\]](#)

[Database Users \[page 707\]](#)

[Add a SAML Identity Provider in SAP HANA Cockpit \[page 733\]](#)

7.2.5.1.8 Assign Roles to a Database User

Roles are the standard mechanism of granting privileges to SAP HANA database users. It is recommended that you assign roles to users instead of granting privileges individually. You can grant roles to users on the [Assign Roles](#) page of the SAP HANA cockpit.

Prerequisites

- The roles you want to assign are available.
Roles in the SAP HANA database can exist as runtime objects only (catalog roles), or in the repository of the SAP HANA database as design-time objects that become runtime objects on activation (repository roles). It is recommended that you model roles as design-time objects. For more information about roles, see *Roles*. For more information about creating repository roles, see *Create a Design-Time Role* in the *SAP HANA Developer Guide (For Web Workbench)*.
- You have the privileges required to grant roles.
If you have the system privilege `ROLE ADMIN`, you can grant any role. Otherwise, the following applies:
 - To grant a catalog role, that is a role created in runtime using SQL, you need to have the role being granted yourself and be authorized to grant it to other users and roles, or the system privilege `ROLE ADMIN`.

- To grant a repository role, that is a role created in the repository of the SAP HANA database, you need the object privilege `EXECUTE` on the procedure `GRANT_ACTIVATED_ROLE`.
- To grant a HDI role, that is a schema-specific role created using the SAP Web IDE and deployed using SAP HANA deployment infrastructure, you need privileges to execute `GRANT_CONTAINER_SCHEMA_ROLES` in the container's API schema, or, if you are a container group administrator, privileges to execute `GRANT_CONTAINER_SCHEMA_ROLES` in the container group's API schema. For more information, see *Prerequisites for Granting and Revoking Privileges and Roles*.

Procedure

1. On the *Overview* page, choose the quicklink *Assign roles to users*.

The *Assign Roles* page opens.

2. Find the user you want to edit.

Detailed information about the user is displayed, including all roles already assigned and who assigned them.

3. Open the user for editing by clicking *Edit*.
4. Grant the user roles by clicking *Assign Roles*, selecting the relevant roles, and clicking *OK*.
5. Optional: Authorize the user to grant the role to other users and roles, by selecting *Grantable to Others*.

i Note

This option is not available for roles created in the repository.

6. Save the user.

The user is granted the selected roles.

Related Information

[Database Roles \[page 759\]](#)

[Prerequisites for Granting and Revoking Privileges and Roles \[page 788\]](#)

7.2.5.1.9 Assign Privileges to a User

It is recommended that you assign roles to users instead of granting privileges individually. However, you can still grant privileges directly to users using the [Assign Privileges](#) app.

Prerequisites

- You have the system privilege USER ADMIN.
- You have the privileges required to grant specific privileges to the user.
To grant SQL privileges, you must have the privilege and/or role yourself and be authorized to grant it to someone else. To grant privileges on activated repository objects, you must be authorized to execute certain stored procedures. For more information, see *Prerequisites for Granting and Revoking Privileges and Roles*.

Procedure

1. On the [Overview](#) page, choose the quicklink [Assign privileges to users](#).

The [Assign Privileges](#) app opens.

2. Find the user you want to edit.

The user's existing privileges are displayed, as well as who assigned them.

3. Assign the required privileges to the user:
 - a. For the relevant privilege type, choose [Edit](#).
 - b. Choose [Add](#) and select the privileges you want to assign.

i Note

For object and package privileges, you must first add the object or package and then add the required privilege to the object or package.

- c. If you want users who have the new role to be able to grant the assigned privilege on to others, choose [Grantable to Others](#).
- d. Save the privilege assignment.
- e. Repeat for further privilege types.

Related Information

[Prerequisites for Granting and Revoking Privileges and Roles \[page 788\]](#)

[Privileges \[page 740\]](#)

[System Privileges \(Reference\) \[page 743\]](#)

[Object Privileges \(Reference\) \[page 749\]](#)

7.2.5.1.10 Change a Database User

You can change an existing database user on the [User](#) page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege `USER ADMIN`.
- If you are integrating SAP HANA database users into a single sign-on (SSO) environment using one or more of the supported mechanisms, the necessary infrastructure must be in place and configured. For more information about single sign-on integration, see the *SAP HANA Security Guide*.

Procedure

1. On the [Overview](#) page, choose the [Manage users](#) link.

The [User](#) page opens. All existing database users are displayed in list format on the left.

2. Find the user you want to change.

→ Tip

Search for a user by entering the name or part of the name in the search box, or create a filter by clicking the [Filter](#) button.

3. Open the user for editing by clicking [Edit](#).
4. Make the required changes.

For more information about the individual fields and settings, see [Database User Details](#).

→ Remember

To change the user's authorization, open the [Assign Roles](#) or [Assign Privileges](#) pages.

Related Information

[Database User Details \[page 777\]](#)

[Assign Roles to a Database User \[page 799\]](#)

[Assign Privileges to a User \[page 801\]](#)

7.2.5.1.11 Deactivate a Database User

Users can be automatically deactivated for security reasons, for example, if they violate password policy rules. However, as a user administrator, you may need to explicitly deactivate a user, for example, if an employee temporarily leaves the company or a security violation is detected. You can deactivate a user on the [User](#) page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege `USER ADMIN`.

Procedure

→ Tip

As an administrator you may want to temporarily deactivate all users in a system except certain administrative users so that these users can perform administration or maintenance tasks. For more information about how to do this without deactivating users individually as described here, see SAP Note 1986645.

1. On the [Overview](#) page, choose the [Manage users](#) link.

The [User](#) page opens. All existing database users are displayed in list format on the left.

2. Find the user you want to deactivate.

→ Tip

Search for a user by entering the name or part of the name in the search box, or create a filter by clicking the [Filter](#) button.

3. Click [Deactivate](#) in the footer bar.

Results

The database user is now deactivated and remains so until you reactivate. The user still exists in the database, but cannot connect to the database any more.

i Note

It may still appear as though deactivated users are still active in the system (for example when a procedure that was created by the user with `DEFINER MODE` is called).

You can activate the user again by clicking [Activate](#) in the footer.

Related Information

[SAP Note 1986645](#)

7.2.5.1.12 Delete a Database User

You may need to delete a database user if an employee leaves your organization for example. You can delete a user with on the *User* page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege `USER ADMIN`.

Procedure

1. On the *Overview* page, choose the *Manage users* link.

The *User* page opens. All existing database users are displayed in list format on the left.

2. Find the user you want to delete.

→ Tip

Search for a user by entering the name or part of the name in the search box, or create a filter by clicking the *Filter* button.

3. Specify whether or not you want to delete dependent objects, such as schemas, tables, views, and procedures with the user.

⚠ Caution

If you choose the *Cascade* option, all objects owned by the user are deleted, and privileges granted to others by the user are revoked. Furthermore, all objects in the user's schema are deleted even if they are owned by a different user. All privileges on these objects are also revoked.

Results

The user is deleted.

7.2.5.1.13 Add a SAML Identity Provider in SAP HANA Cockpit

If you are implementing Security Assertion Markup Language (SAML) to authenticate users accessing SAP HANA via the SQL interface directly (that is using JDBC and ODBC clients), you must add the SAML identity providers for the required users. You can do this using the SAP HANA cockpit.

Prerequisites

- You have created a certificate collection with the purpose **SAML** in the database and have imported the X.509 certificates that will be used to sign the SAML assertions from the identity provider. Ensure that the entire certificate chain of the X.509 certificate is available.

⚠ Caution

We recommend creating certificate collections for individual purposes in the database directly, rather than using trust stores (PSE) in the file system. By default, the same PSE in the file system is shared by all databases for all external communication channels (including HTTP) and certificate-based authentication. Different PSEs must be explicitly configured for tenant databases.

- You have the system privilege USER ADMIN.

Procedure

i Note

While you can configure SAML providers for ODBC/JDBC-based SAML authentication using the SAP HANA cockpit, SAP HANA studio or SQL, always use the SAP HANA XS Administration Tool to configure SAML providers that will be used for **HTTP access via the XS classic server**.

1. In the SAP HANA cockpit, navigate to the [Overview](#) page and choose the [Manage SAML providers](#) link.
2. Add a new identity provider.
 - a. Enter the name of the identity provider.

The following naming conventions apply:

 - Spaces and special characters except underscore (_) are not permitted.
 - The name must start with a letter.
 - The name cannot exceed 127 characters.
 - b. Enter the entity ID.
 - c. Select the appropriate X.509 certificate.

i Note

It is not possible to enter the issuer and subject distinguished names (DNs) manually. If the certificate is not available, click [Go to Certificate Store](#) and import it. Then, return to the [SAML Identity Provider](#) page and start again. For more information, see the section in importing a trusted certificate into the certificate store.

- d. Click *Add*.

Results

The identity provider is now available for mapping to individual database users. You can do this when you create the database user. Alternatively, if the database user already exists, you can change their authentication details.

Related Information

[Import a Trusted Certificate into the Certificate Store \[page 909\]](#)

[Create a Database User \[page 792\]](#)

[Change a Database User \[page 802\]](#)

[Maintaining SAML Providers \(HTTP Access via XS Classic Server\) \[page 1559\]](#)

7.2.5.2 Provisioning Users in SAP HANA Studio

You can use the *User* and *Role* editors of the SAP HANA studio to perform user-provisioning tasks.

[Create a Role in Runtime \[page 807\]](#)

You can create a new role directly in runtime and grant it the privileges and roles necessary for the task or function that it represents. You can create a role with the *Role* editor of the SAP HANA studio.

[Create and Authorize a User \[page 808\]](#)

You create a standard database user for every user who wants to work with the SAP HANA database. When you create a user, you can also configure how the user will be authenticated, as well as which roles and privileges they need. You can create a user with the *User* editor of the SAP HANA studio.

[Create and Authorize a Restricted User \[page 810\]](#)

You create a restricted user for users who access SAP HANA through client applications – full SQL access via an SQL console is not intended. When you create a restricted user, you can also configure how the user will be authenticated, as well as which roles and privileges they need. You can create a restricted user with the *User* editor of the SAP HANA studio.

[Copy a User Based on SAP HANA Repository Roles \[page 813\]](#)

If you are implementing user authorization through roles created in the SAP HANA repository, it is possible to create a new user by copying an existing user. The repository roles granted to the existing user are automatically granted to the new user. SQL roles and individual privileges are **not** granted. You can copy a role in this way using the SAP HANA studio.

[Change a User \[page 814\]](#)

You can change a user's authentication information, grant them new privileges and roles, as well as revoke previously granted privileges and roles. You can change a user with the *User* editor of the SAP HANA studio.

[Delete a User \[page 815\]](#)

You may need to delete a database user if an employee leaves your organization for example. You can delete a user with the *User* editor of the SAP HANA studio.

[Deactivate a User \[page 816\]](#)

Users can be automatically deactivated for security reasons, for example, if they violate password policy rules. However, as a user administrator, you may need to explicitly deactivate a user, for example, if an employee temporarily leaves the company or a security violation is detected. You can deactivate a user with the *User* editor of the SAP HANA studio.

[Reactivate a User \[page 817\]](#)

As a user administrator, you may need to reactivate a user, for example, you explicitly deactivated the user or the user has made too many invalid log-on attempts. You can reactivate a user with the *User* editor of the SAP HANA studio.

[Additional User Parameters \[page 817\]](#)

Several additional user parameters allow you to add more information about a user, for example their e-mail address.

7.2.5.2.1 Create a Role in Runtime

You can create a new role directly in runtime and grant it the privileges and roles necessary for the task or function that it represents. You can create a role with the *Role* editor of the SAP HANA studio.

Prerequisites

- You have the system privilege ROLE ADMIN.
- You have the privileges required to grant privileges and roles to the new role.

Procedure

→ Recommendation

Creating roles in the repository of the SAP HANA database offers more flexibility than creating them in runtime as described here. The recommended approach is therefore to create roles as repository objects. For more information about roles as repository and how to model roles in design time, see the *SAP HANA Developer Guide (For SAP HANA Web Workbench)*.

1. Create a new role:
 - a. In the *Systems* view, choose **► Security ► Roles**.
 - b. From the context menu, choose *New Role*.

The *New Role* editor opens.

2. Specify a unique role name.

The role name can contain all characters, except double quotation marks ("...").

- Optional: Assign the role a runtime namespace by choosing the schema in which to create the role. Role namespaces allow you to reuse roles in different contexts. If you do not select a schema, the role will be created as a global role.

Caution

A role with a namespace will be deleted if the schema is deleted.

- Grant the required roles and privileges.

To authorize a user who has been granted the role to pass on granted roles and privileges to other users, you can select *Grantable to other users and roles*. Note that this option is not available when granting the following:

- Roles created in the repository
- Privileges on objects created in the repository

- Save the role by choosing the  (*Deploy*) button to create the role.

Results

The role is created and appears in the  *Security* > *Roles* folder. It is automatically granted to your user.

Related Information

[Prerequisites for Granting and Revoking Privileges and Roles \[page 788\]](#)

7.2.5.2.2 Create and Authorize a User

You create a standard database user for every user who wants to work with the SAP HANA database. When you create a user, you can also configure how the user will be authenticated, as well as which roles and privileges they need. You can create a user with the *User* editor of the SAP HANA studio.

Prerequisites

- You have the system privilege USER ADMIN.
- You have the privileges required to grant specific privileges and roles to the new user. To grant SQL privileges and roles, you must have the privilege and/or role yourself and be authorized to grant it to others. To grant privileges on activated repository objects, you must be authorized to execute certain stored procedures. For more information, see *Prerequisites for Granting and Revoking Privileges and Roles*.
- If you are integrating SAP HANA database users into a single-sign on (SSO) environment using one or more of the supported mechanisms, the necessary infrastructure must be in place and configured.

Procedure

1. Create the user:
 - a. In the *Systems* view, choose ► *Security* ► *Users* ▾.
 - b. From the context menu, choose *New User*.
The *New User* editor opens on the *User* tab.
2. Specify the new user name.
You must give the user a unique name. User names can contain any CESU-8 characters except for a small subset. For more information, see *Unpermitted Characters in User Names*.
3. Optional: Prevent the user from being able to connect to the database via ODBC and JDBC clients by selecting the corresponding checkbox.

i Note

By default, standard users have access via ODBC and JDBC clients. If ODBC/JDBC client access is disabled, the user can still connect via HTTP. Furthermore, disabling ODBC/JDBC access does not affect the user's authorizations or prevent the user from executing SQL commands via channels other than JDBC/ODBC.

4. Specify the user's properties:

Option	Description
Authenti- cation	<p>You can set up one or more of the following types of user authentication:</p> <ul style="list-style-type: none"> ○ User name/password authentication by a password You can override the password policy setting (<i>force_first_password_change</i>) that forces users to change a password set by a user administrator the first time they log on. This is useful for technical users, for example. ○ Kerberos authentication (external) by specifying the user principal name (UPN) specified in the Microsoft Active Directory or the Kerberos Key Distribution Center as the external ID ○ SAML authentication (external) by selecting the identity provider and then entering the user ID known by the SAML identity provider Alternatively, you can allow the identify provider to map users to the database user by selecting the checkbox in the <i>Any</i> column. ○ X.509 certificates by adding the user's public key certificate(s)
	<div style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 5px;"> <p>i Note</p> <p>X.509 certificates are supported only for HTTP access through SAP HANA XS</p> </div>
	<ul style="list-style-type: none"> ○ SAP logon and assertion tickets
Valid From/ Until	<p>You can specify a validity period for the user. For example, if you are creating a user for a new employee, you can enter their start date in the <i>Valid From</i> field.</p> <p>If you do not enter any values, the user is immediately and indefinitely valid.</p>
Session client	<p>When you create SAP HANA information models (attribute views, analytic views, and calculation views), it is possible to filter the data according to the client specified in table fields such as MANDT or CLIENT. You can specify the client relevant for the user here.</p>

5. Authorize the user by granting the required roles and privileges.

To authorize the user to pass on granted roles and privileges to other users, you can select *Grantable to other users and roles*. Note that this option is not available when granting the following:

- Roles created in the repository
 - Privileges on objects created in the repository
 - Privileges granted on your user
6. Optional: Specify additional user information:
- a. Choose the *User Parameters* tab.
 - b. Add the required parameters.
For more information about the available parameters, see *Custom User Properties*.
7. Save the user by choosing the  (*Deploy*) button.
The system verifies that the user's password complies with the configured password policy and that it is not on the password blacklist.

Results

The user is created and appears in the *Users* folder. A new schema is also created for the user in the catalog. It has the same name as the user. The standard role PUBLIC is always and irrevocably granted. This role allows the user read-only access to system views.

Related Information

[Prerequisites for Granting and Revoking Privileges and Roles \[page 788\]](#)

[Database Users \[page 707\]](#)

[Configure the Database Password Policy and Blacklist in SAP HANA Studio \[page 722\]](#)

[Configure Kerberos for SAP HANA Database Hosts \[page 730\]](#)

[Add a SAML Identity Provider in SAP HANA Studio \[page 732\]](#)

[Maintaining Single Sign-On for SAP HANA XS Applications \[page 1583\]](#)

[Memory Usage in the SAP HANA Database \[page 470\]](#)

[Additional User Parameters \[page 817\]](#)

7.2.5.2.3 Create and Authorize a Restricted User

You create a restricted user for users who access SAP HANA through client applications – full SQL access via an SQL console is not intended. When you create a restricted user, you can also configure how the user will be authenticated, as well as which roles and privileges they need. You can create a restricted user with the *User* editor of the SAP HANA studio.

Prerequisites

- You have the system privilege USER ADMIN.
- You have the privileges required to grant specific privileges and roles to the new user.
To grant SQL privileges and roles, you must have the privilege and/or role yourself and be authorized to grant it to others. To grant privileges on activated repository objects, you must be authorized to execute certain stored procedures. For more information, see *Prerequisites for Granting and Revoking Privileges and Roles*.
- If you are integrating SAP HANA database users into a single-sign on (SSO) environment using one or more of the supported mechanisms, the necessary infrastructure must be in place and configured.

Procedure

1. Create the user:
 - a. In the *Systems* view, choose ► *Security* ► *Users* ▾.
 - b. From the context menu, choose *New Restricted User*.
The *New Restricted User* editor opens.
2. Specify the new user name.
You must give the user a unique name. User names can contain any CESU-8 characters except for a small subset. For more information, see *Unpermitted Characters in User Names*.
3. Optional: Allow the user to connect to the database via ODBC and JDBC clients by deselecting the corresponding checkbox.

By default, restricted users are only able to connect to the database using HTTP. You must explicitly allow access via ODBC and JDBC clients by changing this setting.

For full access to ODBC or JDBC functionality, you must grant restricted users the standard role RESTRICTED_USER_ODBC_ACCESS or RESTRICTED_USER_JDBC_ACCESS.

4. Specify the user's properties:

Option	Description
Authenti- cation	<p>You can set up one or more of the following types of user authentication:</p> <ul style="list-style-type: none">○ User name/password authentication by a password You can override the password policy setting (<i>force_first_password_change</i>) that forces users to change a password set by a user administrator the first time they log on. This is useful for technical users, for example.○ Kerberos authentication (external) by specifying the user principal name (UPN) specified in the Microsoft Active Directory or the Kerberos Key Distribution Center as the external ID○ SAML authentication (external) by selecting the identity provider and then entering the user ID known by the SAML identity provider Alternatively, you can allow the identify provider to map users to the database user by selecting the checkbox in the <i>Any</i> column.○ X.509 certificates by adding the user's public key certificate(s)

Option	Description
	<p>i Note</p> <p>X.509 certificates are supported only for HTTP access through SAP HANA XS</p> <ul style="list-style-type: none"> ○ SAP logon and assertion tickets
Valid From/Until	<p>You can specify a validity period for the user. For example, if you are creating a user for a new employee, you can enter their start date in the <i>Valid From</i> field.</p> <p>If you do not enter any values, the user is immediately and indefinitely valid.</p>
Session client	<p>When you create SAP HANA information models (attribute views, analytic views, and calculation views), it is possible to filter the data according to the client specified in table fields such as MANDT or CLIENT. You can specify the client relevant for the user here.</p>

5. Authorize the user by granting the required roles and privileges.

To authorize the user to pass on granted roles and privileges to other users, you can select *Grantable to other users and roles*. Note that this option is not available when granting the following:

- Roles created in the repository
- Privileges on objects created in the repository
- Privileges granted on your user

6. Optional: Specify additional user information:

- a. Choose the *User Parameters* tab.
- b. Add the required parameters.

For more information about the available parameters, see *Custom User Properties*.

7. Save the user by choosing the  (*Deploy*) button.

The system verifies that the user's password complies with the configured password policy and that it is not on the password blacklist.

Results

The user is created and appears in the *Users* folder. A new schema is also created for the user in the catalog. It has the same name as the user. However, as a restricted user, the user is not authorized to create objects in this schema. For more information about all restrictions, see *Database Users*.

Related Information

[Prerequisites for Granting and Revoking Privileges and Roles \[page 788\]](#)

[Database Users \[page 707\]](#)

[Configure the Database Password Policy and Blacklist in SAP HANA Studio \[page 722\]](#)

[Configure Kerberos for SAP HANA Database Hosts \[page 730\]](#)

[Add a SAML Identity Provider in SAP HANA Studio \[page 732\]](#)

[Maintaining Single Sign-On for SAP HANA XS Applications \[page 1583\]](#)

[Memory Usage in the SAP HANA Database \[page 470\]](#)

7.2.5.2.4 Copy a User Based on SAP HANA Repository Roles

If you are implementing user authorization through roles created in the SAP HANA repository, it is possible to create a new user by copying an existing user. The repository roles granted to the existing user are automatically granted to the new user. SQL roles and individual privileges are **not** granted. You can copy a role in this way using the SAP HANA studio.

Prerequisites

You have the system privilege USER ADMIN and the object privilege EXECUTE on the GRANT_ACTIVATED_ROLE (SYS_REPO) procedure.

Context

Copying a user allows you to create a new user with the same repository roles as the source user automatically granted.

i Note

Only roles created in the SAP HANA repository are granted. SQL roles, including the standard roles delivered with the SAP HANA database (MONITORING, MODELING, and so on) and individual privileges are **not** granted.

Procedure

1. Copy an existing user:
 - a. In the *Systems* view, choose **Security > Users**.
 - b. Right-click the user to be copied and choose *Copy User*.
The *Copy User* editor opens. The repository roles granted to the source user automatically appear on the *Granted Roles* tab.
2. Enter the required user-specific information, that is, user name and authentication details.
3. Grant any additional roles and privileges required by the user.
4. Create the user by choosing the  (*Deploy*) button to create the user.
The system verifies that the user's password complies with the configured password policy and that it is not on the password blacklist.

Results

The user is created and appears in the [Users](#) folder. A new schema is also created for the user in the catalog. It has the same name as the user.

7.2.5.2.5 Change a User

You can change a user's authentication information, grant them new privileges and roles, as well as revoke previously granted privileges and roles. You can change a user with the [User](#) editor of the SAP HANA studio.

Prerequisites

- You have the system privilege USER ADMIN.

Note

A user can change his or her own password without USER ADMIN.

- You have the privileges required to grant specific privileges and roles to the user.
To grant SQL privileges and roles, you must have the privilege and/or role yourself and be authorized to grant it to others. To grant privileges on activated repository objects, you must be authorized to execute certain stored procedures. For more information, see *Prerequisites for Granting and Revoking Privileges and Roles*.
- If you are integrating SAP HANA database users into a single-sign on (SSO) environment using one or more of the supported mechanisms, the necessary infrastructure must be in place and configured.

Procedure

1. Open the user for editing:
 - a. In the [Systems](#) view, choose [Security](#) > [Users](#).
 - b. Open the relevant user.
2. Make the required changes.

You can change the following:

- Authentication methods supported for the user
- Password for user name/password authentication
- External ID for Kerberos authentication, that is the user principal name (UPN) specified in the Microsoft Active Directory or the Kerberos Key Distribution Center
- Identity provider and external user ID for SAML authentication
- User's public key certificate for X.509 certificate authentication (only supported for HTTP access through SAP HANA Extended Services (SAP HANA XS))

- SAP logon and assertion tickets
- Validity period
- Session client
- Granted roles and privileges
- Whether or not the user is allowed to pass on his or her privileges to other users (*Grantable to other users and roles* option)

i Note

This option is **not** available when granting the following:

- Roles created in the repository
- Privileges on objects created in the repository
- Privileges granted on other users

3. Save the user by choosing the  (*Deploy*) button to save the changes.

If you changed the user's password, the system verifies that it complies with the configured password policy and that it is not on the password blacklist.

Related Information

[Prerequisites for Granting and Revoking Privileges and Roles \[page 788\]](#)

[Configure the Database Password Policy and Blacklist in SAP HANA Studio \[page 722\]](#)

[Configure Kerberos for SAP HANA Database Hosts \[page 730\]](#)

[Add a SAML Identity Provider in SAP HANA Studio \[page 732\]](#)

[Maintaining Single Sign-On for SAP HANA XS Applications \[page 1583\]](#)

7.2.5.2.6 Delete a User

You may need to delete a database user if an employee leaves your organization for example. You can delete a user with the *User* editor of the SAP HANA studio.

Prerequisites

You have the system privilege USER ADMIN.

Procedure

1. In the *Systems* view, choose  *Security* > *Users* .

2. Right-click the user you want to delete and choose *Delete*.
3. Confirm whether or not it acceptable that dependent objects, such as schemas, tables, views, and procedures, are deleted with the user (*Cascade* option).

Caution

If you choose the *Cascade* option, all objects owned by the user are deleted, and privileges granted to others by the user are revoked. Furthermore, all objects in the user's schema are deleted even if they are owned by a different user. All privileges on these objects are also revoked.

7.2.5.2.7 Deactivate a User

Users can be automatically deactivated for security reasons, for example, if they violate password policy rules. However, as a user administrator, you may need to explicitly deactivate a user, for example, if an employee temporarily leaves the company or a security violation is detected. You can deactivate a user with the *User* editor of the SAP HANA studio.

Prerequisites

You have the system privilege USER ADMIN.

Procedure

→ Tip

As an administrator you may want to temporarily deactivate all users in a system except certain administrative users so that these users can perform administration or maintenance tasks. For more information about how to do this without deactivating users individually as described here, see SAP Note 1986645.

1. In the *Systems* view, choose  *Security* > *Users* and open the user that you want to deactivate.
2. From the editor toolbar, choose  (*Deactivate User...*)

Results

The database user is now deactivated and remains so until you reactivate. The user still exists in the database, but cannot connect to the database any more. The reason (*explicitly deactivated*) and the time of deactivation are displayed in the user's details.

i Note

It may still appear as though deactivated users are still active in the system (for example when a procedure that was created by the user with DEFINER MODE is called).

Related Information

[SAP Note 1986645](#)

7.2.5.2.8 Reactivate a User

As a user administrator, you may need to reactivate a user, for example, you explicitly deactivated the user or the user has made too many invalid log-on attempts. You can reactivate a user with the *User* editor of the SAP HANA studio.

Prerequisites

You have the system privilege USER ADMIN.

Procedure

1. In the *Systems* view, choose **Security > Users** and open the user that you want to reactivate.
2. From the editor toolbar, choose  (*Activate User...*)
The user is now reactivated.

7.2.5.2.9 Additional User Parameters

Several additional user parameters allow you to add more information about a user, for example their e-mail address.

Parameter	Description
EMAIL ADDRESS	The user's e-mail address The e-mail address must be unique to the user.

Parameter	Description
LOCALE	<p>The user's locale</p> <p>When you create SAP HANA information models (attribute views, analytic views, and calculation views), this parameter can be used to translate information according to the user's locale.</p>
PRIORITY	<p>The priority with which the thread scheduler handles statements executed by the user</p> <p>The priority can be in the range 0-9 with 9 representing the highest priority. 5 is the default priority.</p>
STATEMENT MEMORY LIMIT	<p>The maximum memory (in GB) that can be used by a statement executed by the user</p> <p>The properties <code>statement_memory_limit</code> and <code>statement_memory_limit_threshold</code> in the <code>memory_manager</code> section of the <code>global.ini</code> configuration file are used to limit the memory that can be allocated with respect to statement execution.</p> <p><code>statement_memory_limit_threshold</code> indicates what percentage of the global memory allocation limit must be in use before the specific value of <code>statement_memory_limit</code> is applied. If this memory limit is being applied and a statement execution exceeds it, then the statement is aborted.</p> <p>With this user parameter, you can set a user-specific limit that takes precedence over the global statement memory limit.</p> <p>For more information about memory usage, see <i>Monitoring Memory Usage</i>.</p>
TIME_ZONE	<p>The user's timezone</p> <p>The standard database formats for locale and timezone are supported.</p>

Related Information

[Memory Usage in the SAP HANA Database \[page 470\]](#)
[Setting a Memory Limit for SQL Statements \[page 633\]](#)

7.2.5.3 Provisioning Users Using an LDAP Identity Management Server

The Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing directory services. If you use an LDAP-compliant directory server to manage users and their access to resources, you can leverage LDAP group membership to authorize users.

i Note

An LDAP server can also be used for user authentication in SAP HANA. When LDAP authentication is used, the required database user can be created automatically and authorized based on LDAP group membership.

[Configure LDAP Group Authorization \[page 819\]](#)

To implement LDAP group authorization, you map LDAP groups to SAP HANA roles. You then create, configure, and verify a connection to an LDAP server. Finally, you configure SAP HANA users for LDAP group authorization, if necessary.

[LDAP Provider Details \[page 821\]](#)

To set up a connection to an LDAP server, you must create an LDAP provider in the SAP HANA database. Depending on your requirements, you configure the LDAP server to authenticate and/or authorize users. For LDAP-authenticated users, you can also enable the automatic creation of users in SAP HANA.

Related Information

[Configure an LDAP Server Connection for LDAP User Authentication \[page 736\]](#)

7.2.5.3.1 Configure LDAP Group Authorization

To implement LDAP group authorization, you map LDAP groups to SAP HANA roles. You then create, configure, and verify a connection to an LDAP server. Finally, you configure SAP HANA users for LDAP group authorization, if necessary.

Prerequisites

- An LDAP v3 compliant server
- LDAP provider is already configured.
- You have the following system privileges:
 - `ROLE ADMIN`
 - `LDAP ADMIN`
 - `USER ADMIN`
- If communication between SAP HANA and the LDAP server is to be secured using TLS/SSL, a certificate collection with purpose *LDAP* exists in the database and the certificate of the Certificate Authority (CA) that signed the certificate used by the LDAP server has been added.

Caution

If using the LDAP server for user authentication, you must secure communication to protect the transmission of user passwords between SAP HANA and the LDAP server.

- If you require automatic user creation in SAP HANA, LDAP authentication must be an active authentication mechanism in the SAP HANA database. You can verify this by checking the value of the parameter `[authentication] authentication_methods` in the `global.ini` configuration file.

Procedure

1. Map LDAP groups to roles in the SAP HANA database.

You do this in the role definition in SAP HANA using either the SAP HANA cockpit or the `CREATE ROLE` or `ALTER ROLE` SQL statements. You must specify the unique distinguished name (DN) of an LDAP group.

❁ Example

```
CREATE ROLE Securities_DBA LDAP GROUP
`cn=Securities_DBA,OU=Application,OU=Groups,ou=DatabaseAdmins,cn=Users,dc=lar
gebank,dc=com` ;
```

i Note

It is not possible to include an LDAP group in the definition of design-time roles.

2. Create and configure the LDAP provider in the SAP HANA database using the `CREATE LDAP PROVIDER` statement.

For more information, see the section on LDAP provider details.

→ Remember

You can change the configuration of an LDAP provider using the `ALTER LDAP PROVIDER`.

3. Verify the configuration of the LDAP provider.

You do this using the `VALIDATE LDAP PROVIDER` statement. This will help you to identify and rectify any missing or incorrect settings.

❁ Example

Example 1:

This example verifies that based on the current LDAP configuration and role-to-group mappings in SAP HANA, the SAP HANA user `john` will be granted SAP HANA roles:

```
VALIDATE LDAP PROVIDER my_ldap_provider CHECK USER 'john';
```

Example 2:

This example verifies that based on the current LDAP configuration and role-to-group mappings in SAP HANA, the user `julie` will be created in the SAP HANA database and granted SAP HANA roles once successfully authenticated (using LDAP authentication).

```
VALIDATE LDAP PROVIDER my_ldap_provider CHECK USER CREATION FOR LDAP USER
'julie';
```

4. Configure SAP HANA users for LDAP group authorization.

i Note

This step is only necessary if you have not enabled the LDAP provider for automatic user creation

You can configure users for LDAP group authorization at the time of user creation or later by specifying `LDAP` as the authorization mode using either the SAP HANA cockpit or the `CREATE USER` and `ALTER USER` statements.

Example

```
CREATE USER USER1 PASSWORD enSMRia3s4hS AUTHORIZATION LDAP;
```

If you change the authorization mode of an existing user, any roles and privileges already granted to the user are revoked. For more information, see the *SAP HANA Security Guide*.

→ Tip

To see which authorization mode is configured for a user, display the user in the SAP HANA cockpit or refer to the `AUTHORIZATION_MODE` column of the `USERS` system view.

Related Information

[Managing Client Certificates \[page 900\]](#)

[Create a Database User \[page 792\]](#)

[Change a Database User \[page 802\]](#)

[SAP Note 1848999](#)

7.2.5.3.2 LDAP Provider Details

To set up a connection to an LDAP server, you must create an LDAP provider in the SAP HANA database. Depending on your requirements, you configure the LDAP server to authenticate and/or authorize users. For LDAP-authenticated users, you can also enable the automatic creation of users in SAP HANA.

i Note

The following list describes only the information needed to configure the LDAP provider. The exact syntax of the element is described in the *SAP HANA SQL and System Views Reference*.

Configuration Information	SQL Clause	Description/Example
Name	<code><ldap-provider-name></code>	Name of the LDAP server
How the LDAP server is accessed	<code>CREDENTIAL TYPE</code>	Access takes place using an LDAP server user with permissions to perform searches as specified by the user look-up URL. You must specify the distinguished name (DN) and password of this user.

Configuration Information	SQL Clause	Description/Example
The LDAP URL used to locate a requested user entry on the LDAP server	USER LOOKUP URL	<p>Returns a unique user entry on the LDAP server that corresponds to current SAP HANA user</p> <p>The user look-up URL has the following format:</p> <pre>ldap[s]://<hostname>:<port>/<base_dn>?<attributes>?<scope>?<filter></pre> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p>❖ Example</p> <pre>USER LOOKUP URL 'ldaps://myhostname:389/cn=Users,dc=largebank,dc=com??sub?(&(objectClass=user)(sAMAccountName=*))'</pre> </div> <p>The search for a user entry is based on the SAP HANA user name of the current user. The search filter must include a filter condition in the format '<code><attribute>=*</code>', where <code><attribute></code> is an LDAP attribute whose value is matched against the name of the SAP HANA user.</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p>i Note</p> <p>The search filter must contain a single asterisk (*). An error is returned if either no asterisk or more than one is present.</p> </div> <p>The asterisk (*) is replaced by the SAP HANA user name before the LDAP search query is sent to LDAP server.</p> <p>In the above example <code>(&(objectClass=user)(sAMAccountName=*))</code> becomes <code>(&(objectClass=user)(sAMAccountName=<USER_NAME of the current HANA user>))</code></p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p>i Note</p> <p>The user look-up search query must return a single entry, that is a unique user entry corresponding to the current SAP HANA user. An error occurs if more than one entry is returned.</p> </div> <p>The <code><attributes></code> element must remain empty in the user look-up URL. Any <code><attributes></code> specified in the user look-up URL are ignored. SAP HANA internally constructs the <code><attributes></code> list before executing the user look-up URL as follows:</p> <ul style="list-style-type: none"> • If a nested group URL is specified, <code><attributes></code> is replaced with the attribute specified by <code>ATTRIBUTE DN</code> clause.

Configuration Information	SQL Clause	Description/Example
		<ul style="list-style-type: none">If a nested group URL is not specified, <code><attributes></code> is replaced with the attributes specified by <code>ATTRIBUTE DN</code> and <code>ATTRIBUTE MEMBER_OF</code> clauses.

Configuration Information	SQL Clause	Description/Example
Optional: The LDAP URL used to obtain a user's group membership information, including nested groups, from the LDAP server	NESTED GROUP LOOKUP URL	The nested group look-up URL returns the complete list of groups of which the LDAP user identified by the user look-up URL is a member, including groups with indirect membership.

The nested group look-up URL has the following format:

```
ldap[s]://<hostname>:<port>/<base_dn>?<attributes>?<scope>?<filter>
```

❖ Example

```
NESTED GROUP LOOKUP URL `ldaps://myhostname:389/ou=groupsOU,dc=x??sub?(member:1.2.840.113556.1.4.1941:=*)`
```

The asterisk (*) in the search filter is replaced by the user DN before the nested group look-up URL search query is executed.

i Note

The search filter must contain a single asterisk (*). An error is returned if either no asterisk or more than one is present.

SAP HANA obtains the list of groups from the NESTED GROUP LOOKUP URL query in the following manner:

- If no attributes are specified in the nested group look-up URL, the attribute specified by ATTRIBUTE DN clause is used.

❖ Example

```
'ldap://myhostname:389/ou=groupsOU,dc=x??sub?(member:1.2.840.113556.1.4.1941:=*)'
```

In this case, before sending the above query to the LDAP server to obtain list of groups, SAP HANA replaces the empty attribute list with the attribute specified by the ATTRIBUTE DN clause. Also, '*' in the search filter is replaced with the user DN that is obtained by executing USER LOOKUP URL query.

- If more than one attribute is specified in the nested group look-up URL, then SAP HANA only uses the first attribute from the <attributes> list. Other attributes are ignored.

If the nested group look-up URL query returns more than one entry, then the value of the first attribute from each entry is combined to obtain the complete list of groups for the user.

Configuration Information	SQL Clause	Description/Example
		This attribute is optional if the LDAP provider is being used for user authentication only.
The LDAP attribute that provides the distinguished name (DN) of the LDAP user entry	ATTRIBUTE DN	<p>❖ Example</p> <pre>ATTRIBUTE DN 'distinguishedName'</pre>
Optional: The LDAP attribute that provides a list of groups that a user is a member of	ATTRIBUTE MEMBER_OF	<p>❖ Example</p> <pre>ATTRIBUTE MEMBER_OF 'memberOf'</pre> <p>This attribute is not required if a nested group look-up URL is specified. If both <code>ATTRIBUTE MEMBER_OF</code> and <code>NESTED GROUP LOOKUP URL</code> are specified, <code>NESTED GROUP LOOKUP URL</code> takes precedence.</p> <p>This attribute is optional if the LDAP provider is being used for user authentication only.</p>
Whether or not connections between SAP HANA and the LDAP server are secured using TLS/SSL	SSL ON SSL OFF	<p>This setting applies to both LDAP access authentication and queries, including user lookup and nested groups lookup.</p> <p>If TLS/SSL is enabled, the URL begins with "ldap://".</p> <p>→ Remember</p> <p>The root certificate of the LDAP server must be available in a certificate collection with purpose <i>LDAP</i>.</p> <p>i Note</p> <p>TLS/SSL-secured communication uses the SAP cryptographic library, CommonCryptoLib. For more information, see SAP Note 1848999.</p> <p>Connections to the LDAP server can also be secured using the secure LDAP protocol. In this case, the URL begins with <code>ldaps://</code> and TLS/SSL must be switched off.</p> <p>⚠ Caution</p> <p>Whether the LDAP provider is to be used for LDAP group authorization and/or authentication, you must secure communication between SAP HANA and the LDAP server using the TLS/SSL protocol to protect the transmission of user passwords between SAP HANA and the LDAP server.</p>

Configuration Information	SQL Clause	Description/Example
Whether the LDAP provider is to be used for LDAP group authorization and/or authentication	DEFAULT ON DEFAULT OFF	You can create several LDAP providers but only one can be in use at any time.
Whether or not the LDAP provider is enabled	DISABLE PROVIDER ENABLE PROVIDER	Whether or not the LDAP provider is enabled
Optional: Whether or not the LDAP provider can create a database user in SAP HANA if required	ENABLE USER CREATION FOR LDAP [USER TYPE {STANDARD RESTRICTED}]	<p>The LDAP server creates a user in SAP HANA if all of the following are true:</p> <ul style="list-style-type: none"> • LDAP authentication is enabled both in SAP HANA and for the user logging on. • The user is a member of at least one LDAP group mapped to an SAP HANA role. • A database user with this name does not already exist. <p>By default, the user created is a standard user. If you want restricted users to be created, add the option USER TYPE RESTRICTED.</p> <p>The new database user is automatically configured for LDAP authentication and LDAP group authorization (authorization mode LDAP).</p>

⚠ Caution

If the default LDAP provider is disabled, users configured for LDAP authorization and/or authentication will not be able to log on after the configured role reuse duration has expired. For more information about the reuse duration, see the *SAP HANA Security Guide*.

🌐 Example

The following example creates an LDAP provider for obtaining the LDAP group membership of SAP HANA users via a secure connection. The provider is enabled and set as the default for LDAP group authorization.

```
CREATE LDAP PROVIDER my_ldap_provider CREDENTIAL TYPE 'PASSWORD' USING
'user=cn=LookupAccount,cn=Users,dc=largebank,dc=como;password=hUWe8ZTiQyG' USER
LOOKUP URL 'ldap://myhostname:389/cn=Users,dc=largebank,dc=com??sub?
(&(objectClass=user)(sAMAccountName=*))' ATTRIBUTE DN 'distinguishedName'
ATTRIBUTE MEMBER_OF 'memberOf' SSL ON DEFAULT ON ENABLE PROVIDER;
```

7.3 Auditing Activity in the SAP HANA Database

Auditing provides you with visibility on who did what in the SAP HANA database (or tried to do what) and when. This allows you, for example, to log and monitor read access to sensitive data.

[Managing Auditing in the SAP HANA Cockpit \[page 827\]](#)

Use the SAP HANA cockpit to enable auditing, configure audit trail targets, and create audit policies.

[Managing Auditing in the SAP HANA Studio \[page 837\]](#)

Use the *Security* editor of the SAP HANA studio to enable auditing, configure audit trail targets, and create audit policies.

[Audit Trail Targets \[page 843\]](#)

In production systems, SAP HANA supports syslog and database table as audit trail targets.

[Best Practices and Recommendations for Creating Audit Policies \[page 845\]](#)

7.3.1 Managing Auditing in the SAP HANA Cockpit

Use the SAP HANA cockpit to enable auditing, configure audit trail targets, and create audit policies.

[Activate and Configure Auditing \[page 827\]](#)

The auditing feature of the SAP HANA database allows you to monitor and record selected actions performed in your database. To be able to use this feature, it must first be activated for the database. It is then possible to create and activate the required audit policies. You can do this on the *Auditing* page of the SAP HANA cockpit.

[Create an Audit Policy \[page 829\]](#)

An audit policy defines the actions to be audited, as well as the conditions under which the action must be performed to be relevant for auditing. When an action occurs, the policy is triggered and an audit event is written to the audit trail. Audit policies are database specific. You can create audit policies on the *Auditing* page of the SAP HANA cockpit.

[Delete Audit Entries \[page 831\]](#)

If the audit trail target is or was a database table, you can delete old audit entries, for example to avoid the audit table growing indefinitely. You can do this on the *Auditing* page of the SAP HANA cockpit.

[Auditing Details \[page 832\]](#)

On the *Auditing* page of the SAP HANA cockpit you can view the details of all audit policies in the SAP HANA database, as well as the configured audit trail targets.

[Audit Trail View \[page 835\]](#)

For each occurrence of an audited action, one or more audit entries are created and written to the audit trail. If the audit trail target is a database table, you can view the audit trail in the *Auditing* app of the SAP HANA cockpit. Several options are available for configuring the layout.

7.3.1.1 Activate and Configure Auditing

The auditing feature of the SAP HANA database allows you to monitor and record selected actions performed in your database. To be able to use this feature, it must first be activated for the database. It is then possible to create and activate the required audit policies. You can do this on the *Auditing* page of the SAP HANA cockpit.

Prerequisites

You have the system privileges `AUDIT ADMIN` and `INIFILE ADMIN`.

Procedure

1. On the *Overview* page, navigate to the *Security* area and click the *Auditing* tile.

The *Auditing* page opens.

2. Enable auditing.
 - a. Open the *Configuration* tab and choose *Edit* in the footer bar.
 - b. Set the auditing status to *Enable*.

i Note

You can also enable auditing directly on *Auditing* tile with the on/off switch.

3. Optional: Configure the required audit trail targets.

By default, this is possible only in the **system database**.

You can configure multiple audit trail targets: one for the database (*Overall Audit Trail Target*), and optionally one or more for the severity of audited actions, that is the audit level of the corresponding audit entries. If you do not configure a specific target for an audit level, audit entries are written to the overall audit trail target.

Database table is the default audit trail target for tenant databases and *syslog* for the system database.

i Note

If you are configuring auditing in a tenant database, you cannot change the audit trail targets. This is because the underlying system properties (`[auditing configuration] *_audit_trail_type`) are in the configuration change blacklist `multidb.ini`. Audit trails are by default written to an internal database table of the tenant database. Although not recommended, it is possible to change the audit trail target of a tenant database in the following ways:

- The system administrator changes the audit trail targets for individual tenant databases directly by configuring the relevant system property (`[auditing configuration] *_audit_trail_type`) in the `global.ini` file. For more information about the system properties for configuring audit trail targets and the configuration change blacklist in the *SAP HANA Security Guide*.
- The system administrator removes the relevant system property (`[auditing configuration] *_audit_trail_type`) from the configuration change blacklist, thus enabling the tenant database administrator to change the audit trail target.

⚠ Caution

To ensure the privacy of tenant database audit trails, it is recommended that you do **not** change the default audit trail target (internal database table) of tenant databases.

4. Save your configuration.

Results

Auditing is now activated in your database and you can create the required audit policies.

Related Information

[Audit Trail Targets \[page 843\]](#)

[Default Blacklisted System Properties in Tenant Databases \[page 228\]](#)

7.3.1.2 Create an Audit Policy

An audit policy defines the actions to be audited, as well as the conditions under which the action must be performed to be relevant for auditing. When an action occurs, the policy is triggered and an audit event is written to the audit trail. Audit policies are database specific. You can create audit policies on the [Auditing](#) page of the SAP HANA cockpit.

Prerequisites

You have the system privileges `AUDIT ADMIN` and `INIFILE ADMIN`.

Procedure

1. In the SAP HANA cockpit, navigate to the [Security](#) area of the [Overview](#) page and click the [Auditing](#) tile.
2. On the [Audit Policies](#) tab of the [Auditing](#) page, click the [Create Audit Policy](#) button.
3. Enter the policy name.
4. Optional: Indicate whether you want the audit policy to be immediately enabled (default) or initially disabled on creation.
5. Optional: Select the action status.

The action status specifies when the actions in the policy are to be audited. The following values are possible:

Status	Description
All	The action is audited when the SQL statement is both successfully and unsuccessfully executed.
Successful (default)	The action is audited only when the SQL statement is successfully executed.
Unsuccessful	The action is audited only when the SQL statement is unsuccessfully executed.

i Note

An unsuccessful attempt to execute an action means that the user was not authorized to execute the action. If another error occurs (for example, misspellings in user or object names and syntax errors),

the action is generally not audited. In the case of actions that involve data manipulation (that is, INSERT, SELECT, UPDATE, DELETE, and EXECUTE statements), additional errors (for example, invalidated views) are audited.

6. Optional: Select the audit level.

The audit level specifies the severity of the audit entry written to the audit trail when the actions in the policy occur and ranges from INFO to EMERGENCY. The default level is INFO.

7. Optional: Select one or more policy-specific audit trail targets (system database only).

Audit entries triggered by this policy will be written to the specified audit trail target(s). If you do not specify an audit trail target, entries will be written to the audit trail target for the audit level of the policy if configured, or the audit trail target configured for the system.

8. Specify the actions to be audited by clicking the *Add Actions* button and selecting first the relevant type of activity you want to audit and then the specific actions.

i Note

Only actions of the same type can be combined together in the same policy.

Selecting *All Actions* covers not only all actions that can be audited individually but also actions that cannot otherwise be audited. Such a policy is referred to as a firefighter policy and is useful if you want to audit the actions of a particularly privileged user, for example.

⚠ Caution

The actions that are audited are limited to those that take place inside the database engine while it is running. Therefore, database restart and recovery will not be audited.

9. If necessary, specify the target object(s) to be audited by clicking the *Add Objects* button and selecting the relevant objects.

You must specify a target object if the actions to be audited involve data manipulation, for example, the actions SELECT, INSERT, UPDATE, DELETE, and EXECUTE. The actions in the policy will only be audited when they are performed on the specified object or objects.

When specifying target objects, note the following:

- You can only enter schemas, tables, views, procedures, and functions.
- The target object must be valid for **all actions** in the policy.

10. If necessary, specify the user(s) to be audited by clicking the *Add Users* button and selecting the relevant users.

It is possible to specify that the actions in the policy be audited only when performed by a particular user or users (*Users Included in Policy*). Alternatively, you can specify that the actions in the policy be audited when performed by all users **except** a particular user or users (*Users Excluded from Policy*).

The actions in the policy will only be audited when performed by the specified user(s). If you do not specify a user, the actions will be audited regardless of who performs them.

i Note

You **must** specify a user if you chose to audit all actions.

11. Save the new policy.

Results

The new policy appears in the list of audit policies. Unless you configured it otherwise, the new policy is automatically enabled. This means that when an action in the policy occurs under the conditions defined in the policy, an audit entry is created in the audit trail target(s) configured for the policy. If an action event is audited by multiple audit policies and these audit policies have different audit trail targets, the audit entry is written to all trail targets.

You can disable a policy at any time by changing the policy status. It is also possible to delete a policy.

i Note

Audit policies are not owned by the database user who creates them and therefore will not be deleted if the corresponding database user is deleted.

Related Information

[Audit Trail Targets \[page 843\]](#)

7.3.1.3 Delete Audit Entries

If the audit trail target is or was a database table, you can delete old audit entries, for example to avoid the audit table growing indefinitely. You can do this on the [Auditing](#) page of the SAP HANA cockpit.

Prerequisites

- The audit trail target is or was [Database Table](#).
- You have archived the audit entries that you plan to delete.
- You have the system privilege AUDIT OPERATOR.

Context

Delete audit entries in audit table, for example, to manage the size of the table.

The database monitors the size of the table with respect to the memory allocation limit and issues an alert when it reaches defined values (by default 5%, 7%, 9%, and 11% of the allocation limit). This behavior can be configured with check 64.

i Note

If the table has grown so large that there is not enough memory available to delete old entries as described here, you can use the SQL command `ALTER SYSTEM CLEAR AUDIT LOG ALL` to completely empty the

table. However, even if you archived the audit table beforehand (**recommended**), any new entries written between the time of archiving and the time of clearing may be lost.

Procedure

1. On the *Overview* page, navigate to the *Security* area and click the *Auditing* tile.
The *Auditing* page opens.
2. Choose the *Audit Table Log* tab.
3. Choose the *Delete Audit Entries* and select which audit entries you want to be deleted.
 - Those older than a specific number of days
 - Those created before a date
 - All entries
4. Choose *Delete* to delete the specified entries from the audit table.

Related Information

[Configure Alerting Thresholds \[page 348\]](#)

7.3.1.4 Auditing Details

On the *Auditing* page of the SAP HANA cockpit you can view the details of all audit policies in the SAP HANA database, as well as the configured audit trail targets.

Audit Policies

→ Tip

You can refine the list of audit policies by using the filtering options available in the table toolbar. Filter by policy name or audited action by entering the term directly in the search field, or click the  (*Filter Settings*) button and select the required filter options. To clear all filters, click  (*Clear All Filters*).

Field	Description
<i>Policy Name</i>	Audit policy name

Field	Description
<i>Policy Status</i>	<p>Audit policy status</p> <p>A policy can be either <i>Enabled</i> or <i>Disabled</i>.</p>
<i>Audited Actions</i>	<p>The action to be audited</p> <p>An audit policy can specify several related actions to be audited. For a full list of all actions that can be audited, see the documentation for SQL access control statement CREATE AUDIT POLICY in the <i>SAP HANA SQL and Systems View Reference</i>.</p>
<i>Audited Action Status</i>	<p>When the actions in the policy are to be audited:</p> <ul style="list-style-type: none"> • On successful execution • On unsuccessful execution • On both successful and unsuccessful execution
<i>Audit Level</i>	<p>The severity of the audit entry written to the audit trail when the actions in the policy occur</p> <p>The following audit levels are possible</p> <ul style="list-style-type: none"> • Emergency • Critical • Alert • Warning • Info
<i>Users</i>	<p>User(s) included in the audit policy or excluded from the audit policy</p> <p>Actions in the policy are audited when performed by either the specified user(s) or any user except the specified user(s).</p>
<i>Target Object</i>	<p>Audited object(s)</p> <p>The following target object types are possible:</p> <ul style="list-style-type: none"> • Schemas (and all objects contained within) • Tables • Views • Procedures • Sequences

Field	Description
<i>Audit Trail Target</i>	<p>Policy-specific audit trail target(s)</p> <p>If there is no policy-specific audit trail target, audit entries generated by the policy are written to the audit trail target for the audit level of the policy if configured, or the audit trail target configured for the database. The applicable default audit trail target is always indicated in brackets.</p>

i Note

Policy-specific audit trail targets are only possible in the system database.

Configuration

Field	Description
<i>Overall Audit Trail Target</i>	<p>The default audit trail target for the database</p> <p>If you do not configure a specific target for an audit level or a specific target for an audit policy, audit entries are written to this audit trail target.</p>
<i>Target for Audit Level Alert</i>	The audit trail target to which audit entries with audit level <code>ALERT</code> are written
<i>Target for Audit Level Emergency</i>	The audit trail target to which audit entries with audit level <code>EMERGENCY</code> are written
<i>Target for Audit Level Critical</i>	The audit trail target to which audit entries with audit level <code>CRITICAL</code> are written

i Note

By default, it is not possible to configure audit trail targets in tenant databases. The audit trail target is *Database table*. For more information, see the section on audit trail targets.

Audit Trail

If the audit trail target is or was database table, you can view and manage the audit trail here. For more information, see the section on the audit trail view.

Related Information

[Auditing Activity in the SAP HANA Database \[page 826\]](#)

[Audit Trail Targets \[page 843\]](#)

[Audit Trail View \[page 835\]](#)

[Delete Audit Entries \[page 831\]](#)

7.3.1.5 Audit Trail View

For each occurrence of an audited action, one or more audit entries are created and written to the audit trail. If the audit trail target is a database table, you can view the audit trail in the *Auditing* app of the SAP HANA cockpit. Several options are available for configuring the layout.

To view the audit trail, on the *Overview* page, navigate to the *Security* area and click the *Auditing* tile.

The following sections describe the layout of the audit trail and the options for configuring the layout.

- [Audit Trail Columns \[page 835\]](#)
- [Audit Trail View Configuration \[page 837\]](#)

Audit Trail Columns

Default Columns

Field	Description
Time Stamp	Time of event occurrence (in system local time)
Policy Name	Name of the audit policy that was triggered
Level	Severity of audited action
Status	Execution status of the statement
Client Host	Name of the host where the action occurred
User Name	User who performed the action
Statement	Statement that was executed

Additional Columns

Field	Description
Action	Action that was audited and thus triggered the policy
Application User Name	Application user who performed the action

⚠ Caution

Treat this information with caution. It comes from the application and SAP HANA has no way of verifying its authenticity.

Field	Description
Client Host	Name of the client machine
Client IP	IP address of the client application
Client PID	Process ID of the client process
Client Port	Port of the client process
Comment	Additional information about the audited event
	<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>Currently in case of failed logon attempts, the reason for failure appears in this field.</p> </div>
Connection ID	ID of the session in which the statement was executed
File Name	Configuration file name, for example global.ini
Grantable	Indication of whether the privilege or role was granted with or without GRANT/ADMIN OPTION
Grantee	Name of the target user of the action, for example, grantee in a GRANT statement
Grantee Schema	Name of the schema of a granted or revoked role
Key	Configuration parameter, for example global_auditing_state
Object Name	Name of the object on which an action was performed, for example, a privilege was granted
Original Database	Name of tenant database in which the query originated; relevant for cross-database queries between tenant databases
Original User	Name of database user who executed the query in the origin tenant database; relevant for cross-database queries between tenant databases
Port	Port number
Previous Value	Previous value of the parameter, for example CSVTEXTFILE
Privilege Name	Name of the privilege that was granted or revoked
Role Name	Name of the role that was granted or revoked
Role Schema Name	Name of the schema in which a role was created/dropped, or the schema of a granted/ revoked role
Schema Name	Name of the schema where the action occurred, for example, a privilege was granted on a schema, or a statement was executed on object in a schema
Section	Configuration section name, for example auditing configuration
Service	Name of the service where the action occurred
Value	New parameter value, for example CSTABLE
XS Application User Name	XS application user name, for example, XSA_ADMIN

Audit Trail View Configuration

You can configure the audit trail view by clicking the  (*Settings*) icon. The following options are available.

i Note

Your preferences are not saved when you log out of the cockpit.

Option	Description
<i>Columns</i>	Select the columns you want to see and the order in which they are displayed.
<i>Sort</i>	Sort the audit trail by one or more columns in ascending or descending order.
<i>Filter</i>	Filter the audit trail by creating complex include and exclude filters. <div data-bbox="622 788 708 822" data-label="Section-Header"><h3>i Note</h3></div> <div data-bbox="622 837 876 866" data-label="Text"><p>Values are case sensitive.</p></div>
<i>Group</i>	Organize the audit trail by grouping events according to a particular field (for example, audited action, policy and so on).

Related Information

[Delete Audit Entries \[page 831\]](#)

7.3.2 Managing Auditing in the SAP HANA Studio

Use the *Security* editor of the SAP HANA studio to enable auditing, configure audit trail targets, and create audit policies.

[Activate and Configure Auditing \[page 838\]](#)

The auditing feature of the SAP HANA database allows you to monitor and record selected actions performed in your database. To be able to use this feature, it must first be activated for the database. It is then possible to create and activate the required audit policies. You can do this using the *Security* editor of the SAP HANA studio.

[Create an Audit Policy \[page 839\]](#)

An audit policy defines the actions to be audited, as well as the conditions under which the action must be performed to be relevant for auditing. When an action occurs, the policy is triggered and an audit event is written to the audit trail. Audit policies are database specific. You can create audit policies using the *Security* editor of the SAP HANA studio.

[Delete Audit Entries from the Audit Trail \[page 842\]](#)

If the audit trail target is or was a database table, you can delete old audit entries, for example to avoid the audit table growing indefinitely. You can do this using the *Security* editor of the SAP HANA studio.

7.3.2.1 Activate and Configure Auditing

The auditing feature of the SAP HANA database allows you to monitor and record selected actions performed in your database. To be able to use this feature, it must first be activated for the database. It is then possible to create and activate the required audit policies. You can do this using the *Security* editor of the SAP HANA studio.

Prerequisites

You have the system privilege AUDIT ADMIN.

Procedure

1. Enable auditing:
 - a. In the Security editor of the system to be audited, choose the *Auditing* tab.
 - b. In the *System Settings for Auditing* area, set the global auditing status to *Enabled*.

i Note

In multiple-container systems, you enable auditing for the system database and each tenant database independently.

2. Configure the required audit trail targets.

You can configure multiple audit trail targets: one for the system, and optionally one or more for the severity of audited actions, that is the audit level of the corresponding audit entries. If you do not configure a specific target for an audit level, audit entries are written to the audit trail target configured for the system. For more information about the supported audit trail targets, see *Audit Trail Targets*.

i Note

If you are configuring auditing in a tenant database, you cannot change the audit trail targets. Audit trails are by default written to the internal database table. A system administrator may change the audit trail targets for tenant databases by changing the relevant system property (`[auditing configuration] *_audit_trail_type`) in the `global.ini` file. However, this is not recommended. For more information, see *System Properties for Configuring Auditing* in the *SAP HANA Security Guide*.

3. Save your configuration by choosing the  (*Deploy*) button.

Results

Auditing is now activated in your system (or database) and you can create the required audit policies.

Related Information

[Audit Trail Targets \[page 843\]](#)

7.3.2.2 Create an Audit Policy

An audit policy defines the actions to be audited, as well as the conditions under which the action must be performed to be relevant for auditing. When an action occurs, the policy is triggered and an audit event is written to the audit trail. Audit policies are database specific. You can create audit policies using the [Security](#) editor of the SAP HANA studio.

Prerequisites

You have the system privilege AUDIT ADMIN.

Procedure

1. Create a new policy:
 - a. In the Security editor of the system to be audited, choose the [Auditing](#) tab.
 - b. In the [Audit Policies](#) area, choose [Create New Policy](#).

A new line is added to the list of policies.

2. Enter the policy name.
3. Specify the actions to be audited as follows:
 - a. In the [Audited Actions](#) column, choose the **...** button.
The [Edit Actions Audited by policy_name](#) dialog box appears.
 - b. Select the required actions.

Not all actions can be combined together in the same policy. When you select an action, those actions that are not compatible with the selected action become unavailable for selection.

Selecting [All Actions](#) covers not only all other actions that can be audited individually but also actions that cannot otherwise be audited. Such a policy is referred to as a firefighter policy and is useful if you want to audit the actions of a particularly privileged user.

Caution

The actions that are audited are limited to those that take place inside the database engine while it is running. Therefore, system restart and system recovery will not be audited.

- c. Choose [OK](#).
4. Specify the action status.

The action status specifies when the actions in the policy are to be audited. The following values are possible:

Status	Description
SUCCESSFUL	The action is audited only when the SQL statement is successfully executed.
UNSUCCESSFUL	The action is audited only when the SQL statement is unsuccessfully executed.
ALL	The action is audited when the SQL statement is both successfully and unsuccessfully executed.

i Note

An unsuccessful attempt to execute an action means that the user was not authorized to execute the action. If another error occurs (for example, misspellings in user or object names and syntax errors), the action is generally not audited. In the case of actions that involve data manipulation (that is, INSERT, SELECT, UPDATE, DELETE, and EXECUTE statements), additional errors (for example, invalidated views) are audited.

5. Specify the audit level.

The audit level specifies the severity of the audit entry written to the audit trail when the actions in the policy occur.

6. If necessary, specify the user(s) to be audited.

It is possible to specify that the actions in the policy be audited only when performed by a particular user or users. Alternatively, you can specify that the actions in the policy be audited when performed by all users **except** a particular user or users.

Users do not have to exist before they can be named in an audit policy. However, if a specified user does not exist, it cannot be audited by the audit policy. When the user is subsequently created, the audit policy will apply for the user.

The actions in the policy will only be audited when performed by the specified user(s). If you do not specify a user, the actions will be audited regardless of who performs them.

i Note

You **must** specify a user if you chose to audit all auditable actions.

7. If necessary, specify the target object(s) to be audited.

You must specify a target object if the actions to be audited involve data manipulation, for example, the actions SELECT, INSERT, UPDATE, DELETE, and EXECUTE. The actions in the policy will only be audited when they are performed on the specified object or objects.

When specifying target objects, note the following:

- You can only enter schemas, tables, views, procedures, and functions.
- The target object must be valid for **all actions** in the policy.
- An object does not have to exist before it can be named as the target object of an audit policy. However, if the object does not exist, it cannot be audited by the audit policy. When an object with the specified name is subsequently created, the audit policy will apply for the object, assuming it is of a type that can be audited and the audited action applies to that object type. For example, if the audited action is EXECUTE, the subsequently created object must be a procedure.

- Optional: Specify one or more policy-specific audit trail targets.

Audit entries triggered by this policy will be written to the specified audit trail target(s). If you do not specify an audit trail target, entries will be written to the audit trail target for the audit level of the policy if configured, or the audit trail target configured for the system. For more information about the supported audit trail targets, see *Audit Trail Targets*.

i Note

If you are creating the audit policy in a tenant database, you cannot specify policy-specific audit trail targets. The audit trail targets configured for the system or audit level apply, by default internal database table. A system administrator may change the audit trail targets for tenant databases by changing the relevant system property (`[auditing configuration] *_audit_trail_type`) in the `global.ini` file. However, this is not recommended. For more information, see *System Properties for Configuring Auditing* in the *SAP HANA Security Guide*.

- Save the new policy by choosing the  (*Deploy*) button.

Results

The list of audit policies is saved together with the new policy. The new policy is automatically enabled. This means that when an action in the policy occurs under the conditions defined in the policy, an audit entry is created in the audit trail target(s) configured for the policy. If an action event is audited by multiple audit policies and these audit policies have different audit trail targets, the audit entry is written to all trail targets.

You can disable a policy at any time by changing the policy status. It is also possible to delete a policy.

i Note

Audit policies are not owned by the database user who creates them and therefore will not be deleted if the corresponding database user is deleted.

Related Information

[Audit Trail Targets \[page 843\]](#)

[Best Practices and Recommendations for Creating Audit Policies \[page 845\]](#)

[Editors and Views of the SAP HANA Administration Console \[page 116\]](#)

7.3.2.3 Delete Audit Entries from the Audit Trail

If the audit trail target is or was a database table, you can delete old audit entries, for example to avoid the audit table growing indefinitely. You can do this using the *Security* editor of the SAP HANA studio.

Prerequisites

- The audit trail target is or was *Database Table* (CSTABLE).
- You have archived the audit entries that you plan to delete.
- You have the system privilege AUDIT OPERATOR.

Context

To avoid the audit table growing indefinitely, it is possible to delete old audit entries by truncating the table. The system monitors the size of the table with respect to the memory allocation limit and issues an alert when it reaches defined values (by default 5%, 7%, 9%, and 11% of the allocation limit). This behavior can be configured with check 64.

i Note

If the table has grown so large that there is not enough memory available to delete old entries as described here, you can use the SQL command `ALTER SYSTEM CLEAR AUDIT LOG ALL` to completely empty the table. However, even if you archived the audit table beforehand (**recommended**), any new entries written between the time of archiving and the time of clearing may be lost.

Procedure

1. In the Security editor of the relevant system, choose the *Auditing* tab.
2. Choose the  (*Truncate the database table audit trail*) and select the date and time until which you want audit entries to be deleted.

Results

All entries in the table audit trail up until the specified date are deleted.

Related Information

[Configure Alerting Thresholds with SAP HANA Studio \[page 377\]](#)

7.3.3 Audit Trail Targets

In production systems, SAP HANA supports syslog and database table as audit trail targets.

Audit Trail Target	Description
Internal database table	<p>Using an SAP HANA database table as the target for the audit trail makes it possible to query and analyze auditing information quickly. It also provides a secure and tamper-proof storage location. Audit entries are only accessible through the public system views AUDIT_LOG, XSA_AUDIT_LOG, and the union of these two views ALL_AUDIT_LOG. Only SELECT operations can be performed on this view by users with the system privilege AUDIT OPERATOR or AUDIT ADMIN.</p> <p>To avoid the audit table growing indefinitely, it is possible to delete old audit entries by truncating the table. You can do in the SAP HANA cockpit or with the SQL statement ALTER SYSTEM CLEAR AUDIT LOG. The system monitors the size of the table with respect to the overall memory allocation limit of the system and issues an alert when it reaches defined values (by default 5%, 7%, 9%, and 11% of the allocation limit). This behavior can be configured with check 64 ("Total memory usage of table-based audit log"). Only users with the system privilege AUDIT OPERATOR can truncate the audit table.</p>
Logging system of the Linux operating system (syslog)	<p>The syslog is a secure storage location for the audit trail because not even the database administrator can access or change it. There are also numerous storage possibilities for the syslog, including storing it on other systems. In addition, the syslog is the default log daemon in UNIX systems. The syslog therefore provides a high degree of flexibility and security, as well as integration into a larger system landscape. For more information about how to configure syslog, refer to the documentation of your operating system.</p>

⚠ Caution

If the syslog daemon cannot write the audit trail to its destination, you will not be informed. To avoid a situation in which audited actions are occurring but audit entries are not being written to the audit trail, ensure that the syslog is properly configured and that the audit trail target is accessible and has sufficient space available.

Audit Trail Target	Description
SAP HANA kernel trace	<p>The audit log can be written to a kernel trace file (*.ltc) in the trace directory (/usr/sap/<sid>/<instance>/<host>/trace).</p> <p>The kernel trace output is not human-readable. It must be converted into a CSV-formatted files using the command-line tool <code>hdbtracediag</code> and then loaded into relational tables for SQL analysis.</p> <p><code>hdbtracediag</code> is available on the SAP HANA server at /usr/sap/<sid>/HDB<instance>/exe.</p>

Additionally, the option exists to store the audit trail in a CSV text file. This should only be used for test purposes in non-production systems. A separate CSV file is created for every service that executes SQL.

⚠ Caution

You must not use a CSV text file for a production system as it has severe restrictions.

Firstly, it is not sufficiently secure. By default, the file is written to the same directory as trace files (/usr/sap/<sid>/<instance>/<host>/trace). This means that database users with the system privilege DATA ADMIN, CATALOG READ, TRACE ADMIN, or INIFILE ADMIN can access it. In the SAP HANA database explorer, it is listed under *Database Diagnostics Files*, and at operating system level, any user in the SAPSYS group can access it.

Secondly, audit trails are created for each server in a distributed database system. This makes it more difficult to trace audit events that were executed across multiple servers (distributed execution).

Audit Trails for Tenant Databases

By default, tenant database administrators **cannot** configure audit trail targets independently for their database since the underlying system properties are in the default configuration change blacklist (`multidb.ini`). The default target for all audit trails in tenant databases is internal database table. Although not recommended, it is possible to change the audit trail target of a tenant database in the following ways:

- The system administrator changes the audit trail targets for individual tenant databases directly by configuring the relevant system property (`[auditing configuration] *_audit_trail_type`) in the `global.ini` file. For more information about the system properties for configuring audit trail targets and the configuration change blacklist in the *SAP HANA Security Guide*.
- The system administrator removes the relevant system property (`[auditing configuration] *_audit_trail_type`) from the configuration change blacklist, thus enabling the tenant database administrator to change the audit trail target.

⚠ Caution

To ensure the privacy of tenant database audit trails, it is recommended that you do **not** change the default audit trail target (internal database table) of tenant databases.

Related Information

[Default Blacklisted System Properties in Tenant Databases \[page 228\]](#)

[Delete Audit Entries \[page 831\]](#)

[Configure Alerting Thresholds \[page 348\]](#)

7.3.4 Best Practices and Recommendations for Creating Audit Policies

General Best Practices

To reduce the performance impact of auditing, some basic guidelines for creating audit policies apply.

- Create as few audit policies as possible. It's usually better to have one complex policy than several simple ones.

→ Remember

Some audit actions can't be combined in the same policy.

- Use audit actions that combine other actions where possible.

❖ Example

Audit the `GRANT ANY` action instead of the `GRANT PRIVILEGE` and the `GRANT STRUCTURED PRIVILEGE` actions.

- Create audit policies for DML actions only if required. Auditing DML actions impacts performance more than auditing DDL actions.
- Don't create audit policies for actions that are automatically audited, for example `CREATE AUDIT POLICY`. For a list of actions that are always audited, see the section on the default audit policy in the *SAP HANA Security Guide*.
- Don't create audit policies for database-internal tables that are involved in administration actions. Create policies for the administration actions themselves.

❖ Example

`P_USER_PASSWORD` is an internal database tables that cannot be accessed by any user, not even `SYSTEM`. Changes in these tables are carried out by internal mechanisms, and not by DML operations. Don't include these tables in an audit policy. Instead create an audit policy for changes to users (`ALTER USER` action) instead.

- Create a firefighter policy (that is, a policy that audits all actions for a user) only in exceptional circumstances, for example, to check whether a certain user is being used for everyday work or if a support user has been given access to the system. Firefighter policies may create large amounts of audit data and significantly impact performance if they are used for high-load users.

Recommended Audit Policies

Once auditing is active in the database, certain actions are always audited in the internal audit policy `MandatoryAuditPolicy`. In addition, consider the following recommendations.

Audit policies for administrative activities

At a minimum, we recommend that you create audit policies in development and production systems to audit the following additional administrative activities:

- Changes to SAP HANA configuration files (*.ini files). The relevant audit action is `SYSTEM CONFIGURATION CHANGE`.

Sample Code

```
CREATE AUDIT POLICY "configuration changes" AUDITING SUCCESSFUL SYSTEM
CONFIGURATION CHANGE LEVEL WARNING;
ALTER AUDIT POLICY "configuration changes" ENABLE;
```

- Changes to users. The relevant audit actions are:
 - `CREATE USER`
 - `ALTER USER`
 - `DROP USER`

Sample Code

```
CREATE AUDIT POLICY "user administration" AUDITING SUCCESSFUL CREATE USER,
ALTER USER, DROP USER LEVEL INFO;
ALTER AUDIT POLICY "user administration" ENABLE;
```

- Changes to authorization. The relevant audit actions are:
 - `GRANT ANY`
 - `REVOKE ANY`

Sample Code

```
CREATE AUDIT POLICY "authorizations" AUDITING SUCCESSFUL GRANT ANY, REVOKE
ANY LEVEL INFO;
ALTER AUDIT POLICY "authorizations" ENABLE;
```

If design-time roles and authorizations are used, also audit the execution of the grant/revoke of design-time roles and privileges.

Sample Code

```
CREATE AUDIT POLICY "designtime privileges" AUDITING SUCCESSFUL
EXECUTE on _SYS_REPO.GRANT_ACTIVATED_ANALYTICAL_PRIVILEGE,
_SYS_REPO.GRANT_ACTIVATED_ROLE,
_SYS_REPO.GRANT_APPLICATION_PRIVILEGE,
_SYS_REPO.GRANT_PRIVILEGE_ON_ACTIVATED_CONTENT,
_SYS_REPO.GRANT_SCHEMA_PRIVILEGE_ON_ACTIVATED_CONTENT,
_SYS_REPO.REVOKE_ACTIVATED_ANALYTICAL_PRIVILEGE,
_SYS_REPO.REVOKE_ACTIVATED_ROLE,
_SYS_REPO.REVOKE_APPLICATION_PRIVILEGE,
_SYS_REPO.REVOKE_PRIVILEGE_ON_ACTIVATED_CONTENT,
```

```
_SYS_REPO.REVOKE_SCHEMA_PRIVILEGE_ON_ACTIVATED_CONTENT
LEVEL INFO;
ALTER AUDIT POLICY "designtime privileges" ENABLE;
```

Additional policies in production systems

In production systems, additional audit policies are usually required to log further activities as defined by IT policy and to meet governance and legal requirements such as SOX compliance.

We also recommend auditing not only successful events but unsuccessful events by defining the audit action status `ALL`. Knowing about unsuccessful events might be a prerequisite to discovering an attack on your system.

⚠ Caution

SAP HANA audit policies are defined at the database level and cannot cover all requirements for data protection and privacy. The business semantics of data are part of the application definition and implementation. It is therefore the application that "knows", for example, which tables in the database contain sensitive personal data, or how business level objects, such as sales orders, are mapped to technical objects in the database.

7.4 Managing Data Encryption in SAP HANA

SAP HANA supports data-at-rest encryption and application data encryption.

i Note

For more information about encryption of network communication, see the *SAP HANA Security Guide*.

[Server-Side Data Encryption Services \[page 848\]](#)

SAP HANA features encryption services for encrypting data at rest, as well as an internal encryption service available to applications with data encryption requirements. SAP HANA uses the secure store in the file system functionality to protect all encryption root keys. All passwords on the SAP HANA database server are stored securely.

[SAP HANA Client Secure User Store \(hdbuserstore\) \[page 879\]](#)

The secure user store (`hdbuserstore`) is a tool installed with the SAP HANA client. Use it to store connection information to SAP HANA systems securely on the client so that client applications can connect to SAP HANA without users having to enter this information. It is typically used by scripts connecting to SAP HANA.

[Client-Side Data Encryption \[page 880\]](#)

With client-side data encryption, columns that contain sensitive data, such as credit card numbers or social security numbers, are encrypted by using an encryption key accessible only by the client. Client-side encryption makes encryption transparent to applications and column data is encrypted and decrypted on the client-driver, allowing the application to read and write data in cleartext form.

7.4.1 Server-Side Data Encryption Services

SAP HANA features encryption services for encrypting data at rest, as well as an internal encryption service available to applications with data encryption requirements. SAP HANA uses the secure store in the file system functionality to protect all encryption root keys. All passwords on the SAP HANA database server are stored securely.

Passwords

On the SAP HANA database server, all passwords are stored securely:

- Operating system user passwords are protected by the standard operating system mechanism, `/etc/passwd` file.
- All database user passwords are stored in hashed form using the secure hash algorithm SHA-256.
- Credentials required by SAP HANA applications for outbound connections are securely stored in a database-internal credential store. This internal credential store is in turn secured using the internal application encryption service. For example, in an SAP HANA smart data access scenario, credentials required to access a remote source are protected using the internal application encryption service.

Data-at-Rest Encryption

To protect data saved to disk from unauthorized access at operating system level, the SAP HANA database supports data encryption in the persistence layer for the following types of data:

- Data in data volumes
- Redo logs in log volumes

Data and log backups can also be encrypted.

Security-Relevant Application Data

An internal encryption service is used to encrypt sensitive application data. This includes credentials required by SAP HANA for outbound connections, private keys of the SAP HANA server stored in the database, and data in secure stores defined by application developers using the SAP HANA XS classic `$.security.Store` API.

Secure Store in the File System (SSFS)

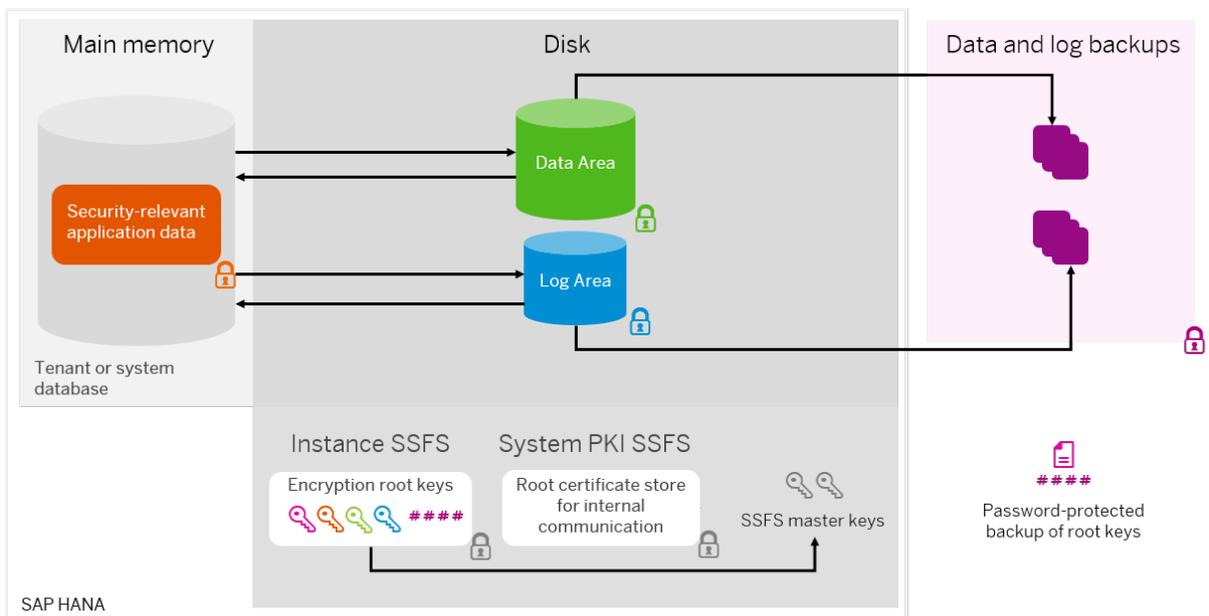
SAP HANA uses two secure stores in the file system: the **instance SSFS** and the **system PKI SSFS**. The instance SSFS protects the root keys used for all data-at-rest encryption services and the internal application encryption service. The system PKI SSFS protects system-internal root certificates required for secure internal communication.

Encryption Services and Keys

The following diagram provides an overview of which data in SAP HANA can be encrypted using a dedicated encryption service, and how all associated encryption root keys are stored in the secure store in the file system of the SAP HANA instance (instance SFSS).

i Note

The following diagram shows only one database. However, a system always has one system database and any number of tenant databases. Every database in the system has its own encryption root keys for each of the available encryption services. The root keys of all databases are stored in the instance SSFS.



Related Information

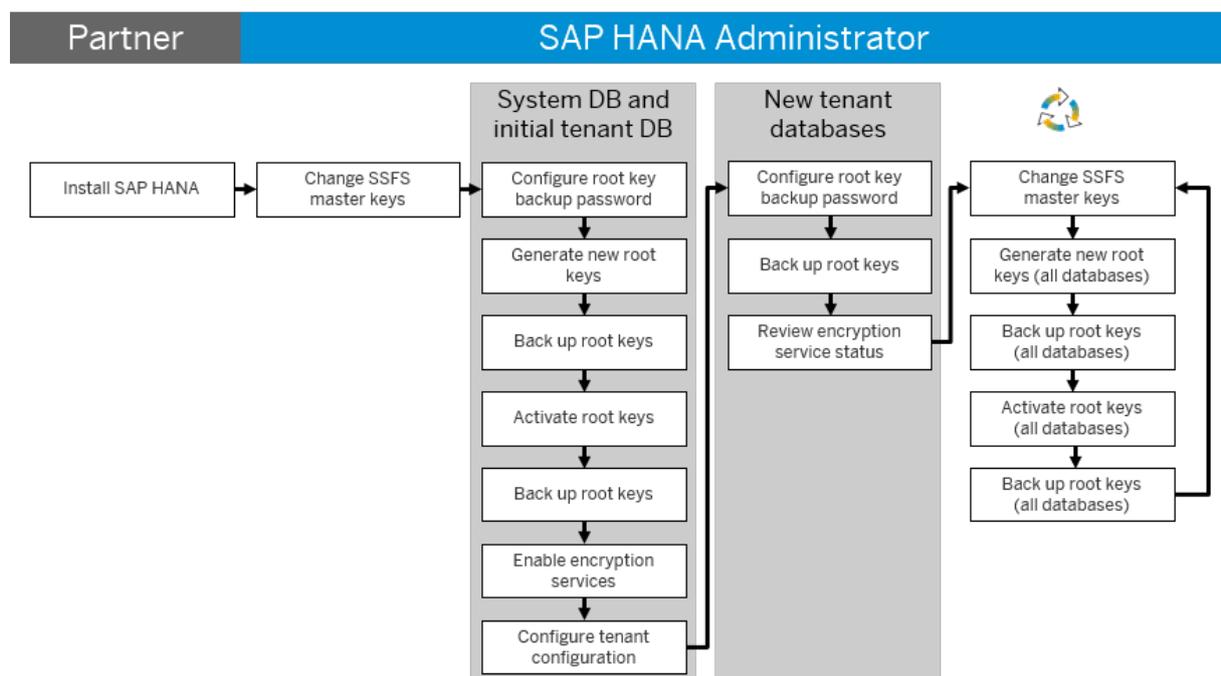
[Encryption Configuration \[page 850\]](#)

7.4.1.1 Encryption Configuration

We recommend that you configure data encryption immediately after handover of your system from a hardware or hosting partner.

First-Time Configuration

The following figure shows the recommended process for configuring encryption in your SAP HANA system for the first time.



Immediately after system handover from your hardware or hosting partner, perform the following steps.

On the SAP HANA server

Change the master keys of the instance SSFS and the system PKI SSFS.

Unique master keys are generated during installation or update. However, if you received your system pre-installed from a hardware or hosting partner, we recommend that you change them immediately after handover to ensure that they are not known outside of your organization. You can also change the master keys any time later.

i Note

In a system-replication configuration, you change the instance SSFS master key on the primary system. To trigger replication of the new key to the secondary system, you must subsequently restart the secondary system. In multi-tier system replication scenarios involving three systems, restart the tier-2 secondary system first, then the tier-3 secondary system. If a secondary system takes over from its replication source before the new master key has been replicated, all systems registered will use the old key from the former secondary system instead.

In the system database

1. Set the password for the root key backup for the system database.
This password is required to securely back up root keys and subsequently restore backed-up root keys during data recovery.

⚠ Caution

The password is stored in the instance SSFS along with the other root keys and used whenever you create a backup of the encryption root keys. The password is required to restore the instance SSFS content before a recovery and should be stored in a separate safe location. Losing this password may result in the database being unrecoverable.

i Note

In a system-replication configuration, set the root key backup password in the primary system only. The password will be propagated to all secondary systems. The secondary systems must be running and replicating.

2. Change the encryption root keys for all encryption services in the system database, that is:
 - Data volume encryption
 - Redo log encryption
 - Data and log backup encryption
 - Internal application encryption

Unique root keys are generated during installation or database creation. However, if you received SAP HANA from a hardware or hosting partner, we recommend that you change them immediately after handover to ensure that they are not known outside of your organization. You can also change root keys any time later.

i Note

In a system-replication configuration, change root keys in the primary system only. New keys will be propagated to all secondary systems. The secondary systems must be running and replicating.

Change all encryption root keys in the system database as follows:

1. Generate new root keys.
2. Back up the new root keys to a root key backup file (*.rkb) in a secure location.

⚠ Caution

Store the root key backup file in a safe location. Losing this file may result in the database being unrecoverable.

3. Activate the new root keys.
4. Back up activated root keys.

You must back up all keys after you generate or activate a key of any type. This ensures that you always have an up-to-date backup of your root keys available for recovery.

For more information about how the key change process works for each of the root key types, see the *SAP HANA Security Guide*.

3. Enable the required encryption services in the system database:
 - Data volume encryption
 - Redo log encryption

- Data and log backup encryption

→ Recommendation

Although SAP HANA provides you with the flexibility to encrypt data volumes, redo logs, and backups independently of each other, if you require full protection in the persistence layer, we recommend that you enable all services.

i Note

It is not necessary to enable the internal application encryption service explicitly. It is available automatically to requesting applications.

i Note

In a system-replication configuration, enable (or disable) encryption in the primary system only. The setting will be propagated to all secondary systems. The secondary systems must be running and replicating.

4. Configure how you want encryption to be handled in new tenant databases.
By default, all encryption services are initially disabled and only tenant database administrators can enable them. You can change this configuration with the following parameters in the `database_initial_encryption` section of the `global.ini` configuration file.
 - `persistence_encryption` (default: off)
 - `log_encryption` (default: off)
 - `backup_encryption` (default: off)
 - `encryption_config_control` (default: local_database)

In the first tenant database (if automatically created during installation)

1. Set the password for the root key backup for the first tenant database.
2. Change the encryption root keys for all encryption services in the first tenant database as described above.
3. Enable the required encryption services in the first tenant database.
Initially, only the tenant database administrator can do this in the tenant database.

i Note

The tenant database administrator can subsequently hand over this configuration control to the system administrator by executing the statement `ALTER SYSTEM ENCRYPTION CONFIGURATION CONTROLLED BY SYSTEM DATABASE.`

→ Remember

In a system-replication configuration, perform all steps in the primary system only. The configuration will be propagated to all secondary systems. The secondary systems must be running and replicating.

In subsequent tenant databases

After you have created an additional tenant database, perform the following steps:

1. Set the password for the root key backup for the tenant database.
2. Back up all root keys to a root key backup file (`*.rkb`) in a secure location.

⚠ Caution

Store the root key backup file in a safe location. Losing this file may result in the database being unrecoverable.

i Note

It is not necessary to change the root keys in new tenant databases. Unique root keys are generated on database creation and cannot be known outside of your organization.

3. Change the status of encryption services in the tenant database if required. Encryption services are initially configured in line with the parameters in the `database_initial_encryption` section of the `global.ini` configuration file as described above. Initially, encryption is disabled. Who can enable or disable encryption services initially depends on how the parameter `encryption_config_control` is configured:
 - If the value of this parameter is `local_database` (default), then only the tenant database administrator can enable or disable encryption from the tenant database.
 - If it is `system_database`, then only the system database administrator can enable or disable encryption from the system database.

i Note

If the tenant database administrator has control over encryption configuration and later wants to hand over this control to the system administrator, the tenant database administrator must execute the statement `ALTER SYSTEM ENCRYPTION CONFIGURATION CONTROLLED BY SYSTEM DATABASE`. If the system administrator has control and wants to hand it over to the tenant database administrator, the system administrator must execute the statement `ALTER DATABASE <database_name> ENCRYPTION CONFIGURATION CONTROLLED BY LOCAL DATABASE`. For simplicity, the system administrator can hand over control to all tenants instead of one by one by executing statement `ALTER SYSTEM ENCRYPTION CONFIGURATION CONTROLLED BY LOCAL DATABASES`.

→ Remember

In a system-replication configuration, perform all steps in the primary system only. The configuration will be propagated to all secondary systems. The secondary systems must be running and replicating.

During operation

Periodically change the SSFS master keys, as well as the encryption root keys in all databases in line with your security policy.

Configuration After Update from a Single-Container System

If you updated from a single-container system, your system has a system database and one tenant database. The existing data encryption configuration is retained. Note the following:

- The SSFS master keys for the system remain unchanged.
- Existing encryption root keys are the encryption root keys of the tenant database. The update process generates new unique root keys for the system database.

- If a root key backup password existed before update, it is the root key backup password of the tenant database. The system database will not have a root key backup password set.
- Encryption services that were enabled before update are enabled in both the system database and the tenant database.

Related Information

[Change the SSFS Master Keys \[page 854\]](#)

[Changing Encryption Root Keys \[page 858\]](#)

[Enabling Encryption of Data and Log Volumes \[page 864\]](#)

7.4.1.2 Change the SSFS Master Keys

The secure stores in the file system (SSFS) used by SAP HANA are protected by unique master keys, generated during installation or update. However, if you received your system pre-installed from a hardware or hosting partner, we recommend that you change these master keys immediately after handover to ensure that they are not known outside your organization.

Prerequisites

- You have shut down the SAP HANA system.
- You have the credentials of the operating system user (<sid>adm) that was created when the system was installed.
- You have the system privilege `INIFILE ADMIN`.

Context

SAP HANA uses the instance SSFS to protect the following encryption root keys:

- The root keys used for:
 - Data volume encryption
 - Redo log encryption
 - Data and log backup encryption
 - Internal application encryption service of the database
- The password of the root key backup
- Encryption configuration information

These root keys protect all encryption keys (and data) used in the SAP HANA database from unauthorized access.

The system database and all tenant databases have their own encryption root keys.

The system PKI SSFS is used to protect the X.509 certificate infrastructure that secures internal SSL/TLS-based communication between hosts in a multiple-host system or between processes of individual databases in the system.

You can change the SSFS master keys using the command line tool `rsecssfx`, which is installed with SAP HANA and available at `/usr/sap/<SID>/HDB<instance>/exe`.

Before changing the SSFS master keys, note the following:

- In a distributed SAP HANA system, every host must be able to access the file location of the instance SSFS master key.
- The SSFS master keys only have to be changed once for the whole instance and not per tenant database.
- In a system-replication configuration, you change the instance SSFS master key on the primary system. To trigger replication of the new key to the secondary system, you must subsequently restart the secondary system. In multi-tier system replication scenarios involving three systems, restart the tier-2 secondary system first, then the tier-3 secondary system. If a secondary system takes over from its replication source before the new master key has been replicated, all systems registered will use the old key from the former secondary system instead.

Procedure

1. Log on to the SAP HANA system host as the operating system user, `<sid>adm`.
2. Change the master key of the **instance SSFS** as follows:
 - a. Re-encrypt the instance SSFS with a new key with the command:

```
export RSEC_SSFS_DATAPATH=/usr/sap/<SID>/SYS/global/hdb/security/ssfs
export RSEC_SSFS_KEYPATH=<path to current key file>
rsecssfx changekey 'rsecssfx generatekey -getPlainValueToConsole'
```

- b. Configure the specified key file location in the `global.ini` configuration file at `/usr/sap/<SID>/SYS/global/hdb/custom/config/global.ini`.

If the file does not exist, create it. Add the following lines:

```
[cryptography]
ssfs_key_file_path = <path to key file>
```

Note

The default path of the key file is `/usr/sap/<sid>/SYS/global/hdb/security/ssfs`. If you change the default path, you may need to reconfigure it in the event of a system rename.

3. Re-encrypt the **system PKI SSFS** with a new key with the following command:

```
export RSEC_SSFS_DATAPATH=/usr/sap/<SID>/SYS/global/security/rsecssfs/data
export RSEC_SSFS_KEYPATH=<path to current key file>
rsecssfx changekey 'rsecssfx generatekey -getPlainValueToConsole'
```

Next Steps

In a system-replication setup, perform the following steps:

1. Configure the location of the instance SSFS master key file on the secondary system(s). The file itself will be automatically copied when you restart the secondary system(s)
2. Restart the secondary system(s) to trigger the replication of the key files.

→ Remember

In multi-tier system replication scenarios involving three systems, restart the tier-2 secondary system first, then the tier-3 secondary system.

For file system-based copies of SAP HANA database installations, you must manually save and restore the instance SSFS master key file. Otherwise data loss can occur.

In regular backup and recovery scenarios, the SSFS must always be restored from the root key backup before a database recovery, unless:

- You have never changed the redo log encryption key.
- You are performing a recovery into the same database from which the backup was taken, and the database's SSFS is intact and contains the latest root key changes.

i Note

It is not necessary to save the system PKI SSFS key file. The system will generate a new system PKI SSFS automatically if required.

Related Information

[Stop a System \[page 183\]](#)

[Configuring SAP HANA System Properties \(INI Files\) \[page 291\]](#)

[Import Backed-up Root Keys \[page 876\]](#)

7.4.1.3 Set the Root Key Backup Password

The root key backup password is required to securely back up the root keys of the database and subsequently to restore the backed-up root keys during data recovery.

Prerequisites

You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.

Procedure

Set the root key backup password.

You can do this using the SAP HANA cockpit or SQL.

Option	Description
SQL	<code>ALTER SYSTEM SET ENCRYPTION ROOT KEYS BACKUP PASSWORD <passphrase></code>
SAP HANA cockpit	<ol style="list-style-type: none">1. In the SAP HANA cockpit, navigate to the <i>Security</i> area of the <i>Overview</i> page.2. Open the <i>Data Encryption Configuration</i> by clicking the <i>Data Encryption</i> tile.3. Choose <i>Manage Keys</i>.4. On the <i>Manage Keys</i> page, click <i>Set Root Key Backup Password</i> and specify the password.

The length and layout of the password must be in line with the database's password policy.

⚠ Caution

If the root key backup already has a password, it will be overwritten.

i Note

In a system-replication configuration, set the root key backup password in the primary system only. The password will be propagated to all secondary systems. The secondary systems must be running and replicating.

Results

The password is set and stored in the secure store in the file system (SSFS) together with other root keys. The password must be set to enable root keys to be backed up securely. You must provide this password to import root keys from the backup into the database before starting a recovery. All root key backups taken after the password is set use this password to protect the backup files.

For more information about root key backup, see the *SAP HANA Security Guide*.

⚠ Caution

The password is stored in the instance SSFS along with the other root keys and used whenever you create a backup of the encryption root keys. The password is required to restore the instance SSFS content before a recovery and should be stored in a separate safe location. Losing this password may result in the database being unrecoverable.

→ Tip

To verify that a password is the same as the one stored in the instance SSFS, use the statement `ALTER SYSTEM VALIDATE ENCRYPTION ROOT KEYS BACKUP PASSWORD <passphrase>`.

Related Information

[Password Policy Configuration Options \[page 723\]](#)

7.4.1.4 Changing Encryption Root Keys

Unique root keys are generated during installation or database creation. However, if you received SAP HANA from a hardware or hosting partner, we recommend that you change them immediately after handover to ensure that they are not known outside of your organization. You can also change root keys any time later.

Change the root keys for the following encryption services immediately after handover of your system and periodically during operation:

- Data volume encryption
- Redo log encryption
- Data and log backup encryption
- Internal application encryption

SAP recommends to always change encryption root keys as follows:

1. Generate new keys.
2. Back up new keys.
3. Activate new keys.
4. Back up activated keys.

You must back up all keys after you generate or activate a key of any type. This ensures that you always have an up-to-date backup of your root keys available for recovery.

For more information about how the key change process works for each of the root key types, see the *SAP HANA Security Guide*.

i Note

In a system-replication configuration, change root keys in the primary system only. New keys will be propagated to all secondary systems. The secondary systems must be running and replicating.

You can change root keys using the SAP HANA cockpit or from the command line.

Related Information

[Change Root Keys Using SAP HANA Cockpit \[page 859\]](#)

[Change Root Keys from the Command Line \[page 860\]](#)

7.4.1.4.1 Change Root Keys Using SAP HANA Cockpit

The process for changing encryption root keys involves first generating the new keys and backing them up, and then activating them and backing them up again. You can change all root keys following this process on the [Manage Keys](#) page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege `ENCRYPTION_ROOT_KEY_ADMIN`.
- You have set the root key backup password.
- The external location to which you plan to back up root keys is accessible.

Procedure

1. In the SAP HANA cockpit, navigate to the [Security](#) area of the [Overview](#) page.
2. Open the [Data Encryption Configuration](#) page by clicking the [Data Encryption](#) tile, and then choose [Manage Keys](#).
3. On the [Manage Keys](#) page, choose [Change Root Keys](#).
4. If you have not already done so, set the root key backup password.

The length and layout of the password must be in line with the database's password policy.

⚠ Caution

The password is stored in the instance SSFS along with the other root keys and used whenever you create a backup of the encryption root keys. The password is required to restore the instance SSFS content before a recovery and should be stored in a separate safe location. Losing this password may result in the database being unrecoverable.

5. Select the root keys that you want to change.
6. Back up all root keys to a secure location.

⚠ Caution

Store the root key backup file in a safe location. Losing this file may result in the database being unrecoverable.

7. Activate the new keys.
8. Back up all root keys again.

Results

If encryption is enabled, new data is encrypted with the new root keys.

On the [Data Encryption](#) page, the active version of changed root keys increments by one and the last changed date is updated.

Related Information

[Set the Root Key Backup Password \[page 856\]](#)

[Password Policy Configuration Options \[page 723\]](#)

7.4.1.4.2 Change Root Keys from the Command Line

The process for changing encryption root keys involves first generating the new keys and backing them up, and then activating them and backing them up again.

Related Information

[Generate New Root Keys \[page 860\]](#)

[Back Up Root Keys \[page 861\]](#)

[Activate New Root Keys \[page 863\]](#)

7.4.1.4.2.1 Generate New Root Keys

The first step in changing encryption root keys is to generate new root keys.

Prerequisites

You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.

Procedure

Generate new root keys for all encryption services by executing the following SQL statements:

Encryption Service	Command
Data volume encryption	ALTER SYSTEM PERSISTENCE ENCRYPTION CREATE NEW ROOT KEY WITHOUT ACTIVATE
Redo log encryption	ALTER SYSTEM LOG ENCRYPTION CREATE NEW ROOT KEY WITHOUT ACTIVATE
Data and log backup encryption	ALTER SYSTEM BACKUP ENCRYPTION CREATE NEW ROOT KEY WITHOUT ACTIVATE
Internal application encryption	ALTER SYSTEM APPLICATION ENCRYPTION CREATE NEW ROOT KEY WITHOUT ACTIVATE

To verify the creation of new root keys, query the system view ENCRYPTION_ROOT_KEYS and check the timestamp of the highest version of all root key types: PERSISTENCE, LOG, BACKUP, and DPAPI. They will have the status PREACTIVE if created with the WITHOUT ACTIVATE option.

7.4.1.4.2 Back Up Root Keys

Once you generate or activate new encryption root keys or create a new tenant database with new root keys, you must back up all root keys.

Prerequisites

- The external location to which you plan to save the backup is accessible.
- You have set the root key backup password using the ALTER SYSTEM SET ENCRYPTION ROOT KEYS BACKUP PASSWORD statement.
- If using hdbnsutil, you have the credentials of the operating system user (<sid>adm), and you know the ID of the database whose root keys you want to back up. You can determine the IDs of all tenant databases by executing the following SQL command in the system database:

```
SELECT DATABASE_NAME,
       CASE WHEN (DBID = '' AND
                 DATABASE_NAME = 'SYSTEMDB')
           THEN 1
           WHEN (DBID = '' AND
                 DATABASE_NAME <> 'SYSTEMDB')
           THEN 3
           ELSE TO_INT(DBID)
           END DATABASE_ID
FROM (SELECT DISTINCT DATABASE_NAME, SUBSTR_AFTER (SUBPATH, '.') AS DBID FROM
      SYS_DATABASES.M_VOLUMES);
```

Procedure

1. Back up the root keys using one of the following methods:

- Using the `hdbnsutil` program
 1. Log on to the SAP HANA server as operating system user `<sid>adm`.
 2. Back up the new keys with the following command:

```
cd /usr/sap/<sid>/<HDBinstance_no>/exe
./hdbnsutil -backupRootKeys <filename>.rkb --dbid=dbid --type='ALL'
```

i Note

- `<dbid>` is the tenant database ID.
- The `<type>` option is the root key type and also accepts the values `PERSISTENCE`, `LOG`, `BACKUP`, and `APPLICATION`. The value `ALL` specifies that root keys of all types are backed up. If you do not specify any value for `<type>`, all root key types are also backed up.

- Extraction

i Note

The database must be online to use this method.

1. Execute the following SQL statement:

```
SELECT ENCRYPTION_ROOT_KEYS_EXTRACT_KEYS ('PERSISTENCE, APPLICATION,
BACKUP, LOG') FROM DUMMY
```

i Note

Execute the statement in the tenant database whose keys are being extracted.

2. Copy the CLOB result and save it to a file at a secure external location. The file must have the extension `*.rkb`.
2. Optional: Validate that you have the password for the root key backup file to ensure that the backup file can be restored:

```
cd /usr/sap/<sid>/<HDBinstance_no>/exe
./hdbnsutil -validateRootKeysBackup <filename> [--password=<passphrase>]
```

→ Recommendation

We recommend that you do not enter the password on the command line. You will be interactively prompted to enter it. In this way, you avoid unintentionally leaving the password in command history and making it visible in process monitoring tools provided by the operating system.

7.4.1.4.2.3 Activate New Root Keys

Activate new encryption root keys so that they can be used to encrypt new data.

Prerequisites

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- You have backed up the new encryption root keys. You can verify whether or not root keys are backed up by querying system view `ENCRYPTION_ROOT_KEYS`.

Procedure

Activate the new root keys by executing the following SQL statements:

Root Key	Command
Data volume encryption	<pre>ALTER SYSTEM PERSISTENCE ENCRYPTION ACTIVATE NEW ROOT KEY</pre>
Redo log encryption	<pre>ALTER SYSTEM LOG ENCRYPTION ACTIVATE NEW ROOT KEY</pre>
Data and log backup encryption	<pre>ALTER SYSTEM BACKUP ENCRYPTION ACTIVATE NEW ROOT KEY</pre>
Internal application encryption	<pre>ALTER SYSTEM APPLICATION ENCRYPTION ACTIVATE NEW ROOT KEY</pre>

If encryption is enabled, new data is encrypted with the new root keys.

i Note

It is not necessary to enable the internal application encryption service explicitly. It is available automatically to requesting applications.

Next Steps

Back up the activated root keys.

Related Information

[Back Up Root Keys \[page 861\]](#)

7.4.1.5 Back Up Root Keys

After you have generated or activated new encryption root keys, or created a new tenant database with new root keys, you must back up all root keys.

Prerequisites

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- You have set the root key backup password.
- The external location to which you plan to back up root keys is accessible.

Procedure

1. In the SAP HANA cockpit, navigate to the *Security* area of the *Overview* page.
2. Open the *Data Encryption Configuration* page by clicking the *Data Encryption* tile.
3. Choose *Back Up Root Keys*.
4. Save the root key backup file to a secure location.

⚠ Caution

Store the root key backup file in a safe location. Losing this file may result in the database being unrecoverable.

7.4.1.6 Enabling Encryption of Data and Log Volumes

You can enable data volume encryption and redo log encryption in a new SAP HANA database or in an existing operational database.

→ Recommendation

Although SAP HANA provides you with the flexibility to encrypt data volumes, redo logs, and backups independently of each other, if you require full protection in the persistence layer, we recommend that you enable all services.

[Enable Data and Log Volume Encryption in a New SAP HANA Database \[page 865\]](#)

The recommended time to enable data and log volume encryption is immediately after tenant database creation. If you received SAP HANA from a hardware or hosting partner, enable encryption after handover.

[Enable Data and Log Volume Encryption in an Existing SAP HANA Database \[page 867\]](#)

There are two ways to enable encryption in an existing operational SAP HANA database. The recommended way involves dropping and re-creating the tenant database. If this is not possible (for

example, because it would result in too much downtime), you can enable encryption immediately. However, be aware that your data will only be fully protected after some delay.

7.4.1.6.1 Enable Data and Log Volume Encryption in a New SAP HANA Database

The recommended time to enable data and log volume encryption is immediately after tenant database creation. If you received SAP HANA from a hardware or hosting partner, enable encryption after handover.

Prerequisites

- You know whether encryption must be enabled/disabled in the tenant database directly or from the system database.

By default, encryption can be enabled/disabled only in the tenant database. To see how your database is configured, query the system view `SYS.M_ENCRYPTION_OVERVIEW`, or from the system database `SYS_DATABASES.M_ENCRYPTION_OVERVIEW`. For more information about how to switch control, see the section on encryption configuration.

i Note

If encryption in the tenant database must be enabled by the system database administrator, the system privilege `DATABASE ADMIN` is required.

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- You have changed and backed up the encryption root keys if necessary. SAP HANA generates unique root keys on installation or database creation. However, if you received SAP HANA from a hardware or hosting partner, you might want to change the root keys used for data volume encryption and redo log encryption to ensure they are not known outside your organization..

i Note

In a system-replication configuration, change root keys in the primary system only. New keys will be propagated to all secondary systems. The secondary systems must be running and replicating.

Procedure

i Note

In a system-replication configuration, enable (or disable) encryption in the primary system only. The setting will be propagated to all secondary systems. The secondary systems must be running and replicating.

1. Enable data volume encryption.

You can do this using SQL, the SAP HANA cockpit, or the SAP HANA studio.

Option	Description
SQL	<ul style="list-style-type: none"> ○ If the tenant database has control: ALTER SYSTEM PERSISTENCE ENCRYPTION ON ○ If the system database has control: ALTER DATABASE <database_name> PERSISTENCE ENCRYPTION ON
SAP HANA cockpit	<p>Tenant database control only:</p> <ol style="list-style-type: none"> 1. In the SAP HANA cockpit, navigate to the <i>Security</i> area of the <i>Overview</i> page. 2. On the <i>Data Storage Security</i> block, enable data volume encryption with the on/off switch.
<p>i Note</p> <p>You can also enable data volume encryption on the <i>Data Volume Encryption</i> page.</p>	
SAP HANA studio	<p>Tenant database control only:</p> <ol style="list-style-type: none"> 1. In the Security editor of the system or database to be encrypted, choose the <i>Data Volume Encryption</i> tab. 2. Select <i>Encrypt data volumes</i> and choose  (<i>Deploy</i>).

2. Enable redo log encryption.

You can do this using SQL or the SAP HANA cockpit.

Option	Description
SQL	<ul style="list-style-type: none"> ○ If the tenant database has control: ALTER SYSTEM LOG ENCRYPTION ON ○ If the system database has control: ALTER DATABASE <database_name> LOG ENCRYPTION ON
SAP HANA cockpit	<p>Tenant database control:</p> <ol style="list-style-type: none"> 1. In the SAP HANA cockpit, navigate to the <i>Security</i> area of the <i>Overview</i> page. 2. On the <i>Data Storage Security</i> block, enable redo log encryption with the on/off switch.
<p>i Note</p> <p>You can also disable redo log encryption on the <i>Data Volume Encryption</i> page.</p>	

Results

All data persisted to data volumes is encrypted and all future redo log entries persisted to log volumes are encrypted.

You can verify the status of data volume encryption and redo log encryption in the system view M_ENCRYPTION_OVERVIEW or on the *Data Encryption* tile of the SAP HANA cockpit.

Related Information

[Encryption Configuration \[page 850\]](#)

7.4.1.6.2 Enable Data and Log Volume Encryption in an Existing SAP HANA Database

There are two ways to enable encryption in an existing operational SAP HANA database. The recommended way involves dropping and re-creating the tenant database. If this is not possible (for example, because it would result in too much downtime), you can enable encryption immediately. However, be aware that your data will only be fully protected after some delay.

Context

i Note

Ideally, you enable encryption in the system database and tenant databases immediately on creation.

Enabling data volume encryption and redo log encryption does not increase data size.

7.4.1.6.2.1 Enable Data and Log Volume Encryption with Database Re-Creation

The recommended way to enable data volume encryption and redo log encryption in an existing operational SAP HANA database is after first dropping and re-creating the tenant database.

Prerequisites

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- You have the system privilege `DATABASE ADMIN`.
- You have the privileges required to perform backup and recovery.

Context

Enabling data volume encryption and redo log encryption after re-creating your tenant database ensures that new encryption root keys are generated. In addition, it provides complete protection. If you enable encryption without re-creating the database, only the pages in use within the data volumes will be encrypted. Pages in data volumes that are not in use may still contain old content and will only be overwritten and encrypted over time.

This means that your data in data volumes will only be fully protected after some delay. In addition, only future redo log entries will be encrypted. Existing redo log files are not encrypted.

For more information about this recommendation, see SAP Note 2159014.

The following is the overall process for enabling encryption with database re-creation. For more information on individual steps, see the related documentation.

Procedure

1. Perform a data backup.
2. Drop the tenant database.
All volumes are removed
3. Create the tenant database again.

i Note

In a system-replication configuration, all secondary systems must be now be running and replicating before the next steps can be performed.

4. Set the password for the root key backup in the tenant database.

i Note

In a system-replication configuration, set the root key backup password in the primary system only. The password will be propagated to all secondary systems. The secondary systems must be running and replicating.

5. Back up all root keys to a root key backup file (*.rkb) in a secure location.
6. Enable data volume and redo log encryption.

By default, encryption is initially disabled in a new database and can be enabled only in the tenant database. This initial configuration is controlled by the parameters in the `database_initial_encryption` section of the `global.ini` configuration file. To see how the database is configured, query the system view `SYS.M_ENCRYPTION_OVERVIEW`, or from the system database `SYS_DATABASES.M_ENCRYPTION_OVERVIEW`.

i Note

In a system-replication configuration, enable (or disable) encryption in the primary system only. The setting will be propagated to all secondary systems. The secondary systems must be running and replicating.

7. Recover your tenant database.

Results

All data persisted to data volumes is encrypted and all future redo log entries persisted to log volumes are encrypted.

You can verify the status of data volume encryption and redo log encryption in the system view `M_ENCRYPTION_OVERVIEW` or on the *Data Encryption* tile of the SAP HANA cockpit.

Related Information

[Encryption Configuration \[page 850\]](#)

[Creating Backups \[page 1313\]](#)

[Delete a Tenant Database \[page 216\]](#)

[Create a Tenant Database \[page 210\]](#)

[Set the Root Key Backup Password \[page 856\]](#)

[Back Up Root Keys \[page 861\]](#)

[Recovering an SAP HANA Database \[page 1347\]](#)

[Enable Data and Log Volume Encryption in a New SAP HANA Database \[page 865\]](#)

[SAP Note 2159014](#)

7.4.1.6.2.2 Enable Data and Log Volume Encryption Without Database Re-Creation

If it is not possible to drop and re-create your SAP HANA database to enable encryption and redo log encryption, for example, because it would result in too much downtime, you can enable encryption immediately. However, this is not recommended because your data will only be fully protected after some delay.

Prerequisites

- You know whether encryption must be enabled/disabled in the tenant database directly or from the system database.

By default, encryption can be enabled/disabled only in the tenant database. To see how your database is configured, query the system view `SYS.M_ENCRYPTION_OVERVIEW`, or from the system database `SYS_DATABASES.M_ENCRYPTION_OVERVIEW`. For more information about how to switch control, see the section on encryption configuration.

i Note

If encryption in the tenant database must be enabled by the system database administrator, the system privilege `DATABASE ADMIN` is required.

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- You have changed and backed up the encryption root keys if necessary. SAP HANA generates unique root keys on installation or database creation. However, if you received SAP HANA from a hardware or hosting partner, you might want to change the root keys used for data volume encryption and redo log encryption to ensure they are not known outside your organization.

i Note

In a system-replication configuration, change the root keys used for data volume encryption and log volume encryption in the primary system only. The new keys will be propagated to all secondary systems.

Context

For maximum protection, we recommend that you drop and re-create your SAP HANA database before enabling data volume encryption and redo log encryption. If you enable encryption once the database has been operational, only the pages in use within the data volumes will be encrypted. Pages in data volumes that are not in use may still contain old content and will only be overwritten and encrypted over time. This means that your data in data volumes will only be fully protected after some delay. In addition, only future redo log entries will be encrypted. Existing redo log files are not encrypted.

For more information about this recommendation, see SAP Note 2159014.

i Note

In a system-replication configuration, enable (or disable) encryption in the primary system only. The setting will be propagated to all secondary systems. The secondary systems must be running and replicating.

Procedure

1. Enable data volume encryption.

You can do this using SQL, the SAP HANA cockpit, or the SAP HANA studio.

Option	Description
SQL	<ul style="list-style-type: none">○ If the tenant database has control: <code>ALTER SYSTEM PERSISTENCE ENCRYPTION ON</code>○ If the system database has control: <code>ALTER DATABASE <database_name> PERSISTENCE ENCRYPTION ON</code>
SAP HANA cockpit	Tenant database control only: <ol style="list-style-type: none">1. In the SAP HANA cockpit, navigate to the Security area of the Overview page.2. On the Data Storage Security block, enable data volume encryption with the on/off switch.
<h3>i Note</h3> <p>You can also enable data volume encryption on the Data Volume Encryption page.</p>	
SAP HANA studio	Tenant database control only: <ol style="list-style-type: none">1. In the Security editor of the system or database to be encrypted, choose the Data Volume Encryption tab.2. Select Encrypt data volumes and choose  (Deploy).

2. Enable redo log encryption.

You can do this using SQL or the SAP HANA cockpit.

Option	Description
SQL	<ul style="list-style-type: none">○ If the tenant database has control: ALTER SYSTEM LOG ENCRYPTION ON○ If the system database has control: ALTER DATABASE <database_name> LOG ENCRYPTION ON
SAP HANA cockpit	Tenant database control: <ol style="list-style-type: none">1. In the SAP HANA cockpit, navigate to the <i>Security</i> area of the <i>Overview</i> page.2. On the <i>Data Storage Security</i> block, enable redo log encryption with the on/off switch.

i Note

You can also disable redo log encryption on the [Data Volume Encryption](#) page.

Results

Encryption is now active for all new data saved to disk as of the next savepoint operation. Existing data starts being encrypted in the background. Only after this process has completed is all your data encrypted.

You can monitor the progress of data volume encryption service by service in the SAP HANA cockpit and SAP HANA studio. Once encryption of a data volume has completed, the status changes to *Encrypted*.

i Note

In the SAP HANA studio, you must refresh (🔄) the editor to see status changes.

→ Remember

Due to the shadow memory nature of SAP HANA database persistence, the data area may still contain outdated, unencrypted versions of pages.

All future redo log entries persisted to log volumes are encrypted. Previous unencrypted redo log entries remain unencrypted.

→ Remember

Existing redo log files will not be encrypted until they are overwritten.

Related Information

[Encryption Configuration \[page 850\]](#)

[Configuring SAP HANA System Replication \[page 1089\]](#)

[SAP Note 2159014](#)

7.4.1.7 Enable Encryption of Data and Log Backups

You can enable encryption of full data backups, delta data backups, and log backups in an SAP HANA database at any time.

Prerequisites

- You know whether encryption must be enabled/disabled in the tenant database directly or from the system database.

By default, encryption can be enabled/disabled only in the tenant database. To see how your database is configured, query the system view `SYS.M_ENCRYPTION_OVERVIEW`, or from the system database `SYS_DATABASES.M_ENCRYPTION_OVERVIEW`. For more information about how to switch control, see the section on encryption configuration.

i Note

If encryption in the tenant database must be enabled by the system database administrator, the system privilege `DATABASE ADMIN` is required.

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- You have changed and backed up the encryption root keys if necessary. SAP HANA generates unique root keys on installation or database creation. However, if you received SAP HANA from a hardware or hosting partner, consider changing the backup encryption root key to ensure it is not known outside your organization.

i Note

In a system-replication configuration, change the root keys used for data volume encryption and log volume encryption in the primary system only. The new keys will be propagated to all secondary systems.

Procedure

Enable backup encryption.

You can do this using SQL or the SAP HANA cockpit.

Option	Description
SQL	<ul style="list-style-type: none">If the tenant database has control: <code>ALTER SYSTEM BACKUP ENCRYPTION ON</code>If the system database has control: <code>ALTER DATABASE <database_name> BACKUP ENCRYPTION ON</code>
SAP HANA cockpit	Tenant database control only: <ol style="list-style-type: none">In the SAP HANA cockpit, navigate to the <i>Security</i> area of the <i>Overview</i> page.

Option	Description
	2. On the <i>Data Storage Security</i> block, enable backup encryption with the on/off switch.
	<p>i Note</p> <p>You can also enable data volume encryption on the <i>Data Volume Encryption</i> page</p>

Results

Backup encryption is enabled. Subsequent log backups, as well as full backups and delta data backups will be encrypted.

i Note

If backup encryption is active, a data snapshot is **not automatically encrypted**. For more information, see *Points to Note: SAP HANA Backup Encryption*.

Related Information

[Encryption Configuration \[page 850\]](#)

[SAP HANA Backup Encryption \[page 1252\]](#)

[Changing Encryption Root Keys \[page 858\]](#)

7.4.1.8 Disable Data Encryption

Disabling data volume encryption triggers the decryption of all encrypted data. Newly persisted data is not encrypted. Disabling redo log encryption makes sure that future redo log entries are not encrypted when they are written to disk.

Prerequisites

- You know whether encryption must be enabled/disabled in the tenant database directly or from the system database.
By default, encryption can be enabled/disabled only in the tenant database. To see how your database is configured, query the system view `SYS.M_ENCRYPTION_OVERVIEW`, or from the system database `SYS_DATABASES.M_ENCRYPTION_OVERVIEW`. For more information about how to switch control, see the section on encryption configuration.

i Note

If encryption in the tenant database must be enabled by the system database administrator, the system privilege `DATABASE ADMIN` is required.

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.

Procedure

i Note

In a system-replication configuration, enable (or disable) encryption in the primary system only. The setting will be propagated to all secondary systems. The secondary systems must be running and replicating.

Disable the required encryption service.

You can do this using SQL or SAP HANA cockpit.

Option	Description
SQL	<p>If the tenant database has control:</p> <ul style="list-style-type: none">◦ Data volume encryption: <code>ALTER SYSTEM PERSISTENCE ENCRYPTION OFF</code>◦ Redo log encryption: <code>ALTER SYSTEM LOG ENCRYPTION OFF</code>◦ Backup encryption: <code>ALTER SYSTEM BACKUP ENCRYPTION OFF</code> <p>If the system database has control:</p> <ul style="list-style-type: none">◦ Data volume encryption: <code>ALTER DATABASE <database_name> PERSISTENCE ENCRYPTION OFF</code>◦ Redo log encryption: <code>ALTER DATABASE <database_name> LOG ENCRYPTION OFF</code>◦ Backup encryption: <code>ALTER DATABASE <database_name> BACKUP ENCRYPTION OFF</code>
SAP HANA cockpit	<p>Tenant database control only:</p> <ol style="list-style-type: none">1. In the SAP HANA cockpit, navigate to the Security area of the Overview page.2. On the Data Encryption tile, disable the relevant encryption service using the on/off switch. <div data-bbox="866 1688 1402 1823"><h3>i Note</h3><p>You can also disable each encryption service on the Data Encryption Configuration page.</p></div>

Results

Data volume encryption

Data starts being decrypted in the background. Depending on the size of the SAP HANA database, this process can be very time consuming. Only after this process has completed is all your data decrypted. Newly persisted data is not encrypted.

You can monitor the progress of data volume decryption service by service. Once decryption of a data volume has completed, the status changes to *Unencrypted*.

Redo log encryption

New redo log entries are not encrypted. Existing redo log entries are not decrypted. Log entries will only be fully unencrypted when all encrypted entries have been overwritten.

Backup encryption

New data backups, delta backups, and log backups are not encrypted. On an unencrypted data volume, data snapshots are also unencrypted.

Related Information

[Encryption Configuration \[page 850\]](#)

7.4.1.9 Change the Page Encryption Key Used for Data Volume Encryption

It is recommended that you periodically change the encryption key used to encrypt pages in the data area in line with your organization's security policy. If necessary, you can re-encrypt the entire data area with the new key. You can change the page encryption key in the SAP HANA studio.

Prerequisites

You have the system privilege `RESOURCE ADMIN` or `ENCRYPTION ROOT KEY ADMIN`.

Context

Changing the encryption key used to encrypt pages in the data area limits the potential impact of a key being compromised. It is recommended that you do so periodically in line with your security policy. You can trigger the creation of a new (randomly generated) page encryption key at any time. This new key is used to encrypt pages as of the next savepoint operation. By default, pages that were previously written to disk are not re-

encrypted. However, you may need or want to re-encrypt your entire data area with the new key. For example, you have a lot of encryption keys in your system, an encryption key was compromised, or your organization's security policy requires that all data be encrypted with keys not older than a certain age.

You can see all encryption keys in your system and their validity periods in the monitoring view `M_PERSISTENCE_ENCRYPTION_KEYS`.

Procedure

1. In the Security editor, choose the *Data Volume Encryption* tab.
2. Choose the  (*Create new page encryption key*) button.

To have the entire data area re-encrypted with the new key in addition, choose *Force all data to be re-encrypted*.

After the next savepoint operation, a new random encryption key is generated. This key will be used to encrypt pages as of the next savepoint operation. Depending on the workload of the database, this may not happen for some time. Pages that were previously written to disk are only re-encrypted if you selected the corresponding option. If this is the case, old pages are first decrypted using the old key and then re-encrypted with the new key.

Results

You can verify the result of a key change in the monitoring view `M_PERSISTENCE_ENCRYPTION_STATUS`. The column `KEY_CHANGE_WITH_NEXT_SAVEPOINT` contains the value `TRUE`.

i Note

Encryption keys that are no longer in use are automatically removed the next time the database is restarted.

7.4.1.10 Import Backed-up Root Keys

Before performing a recovery from encrypted data and log backups, you must import backed-up root keys. The imported keys are then used to initialize the instance SSFS. In this way, the SSFS has the consistent versioned key information required to recover encrypted data backups and replay redo logs.

Prerequisites

- You have the credentials of the operating system user (`<sid>adm`).
- You can log on to the system database and have the system privileges `DATABASE STOP`.

- The location of the root key backup file (*.rkb) is accessible.
- If using `hdbnsutil`, you know the ID of the database whose root keys you want to back up. You can determine the IDs of all tenant databases by executing the following SQL command in the system database:

```
SELECT DATABASE_NAME,
       CASE WHEN (DBID = '' AND
                 DATABASE_NAME = 'SYSTEMDB')
           THEN 1
           WHEN (DBID = '' AND
                 DATABASE_NAME <> 'SYSTEMDB')
           THEN 3
           ELSE TO_INT(DBID)
           END DATABASE_ID
FROM (SELECT DISTINCT DATABASE_NAME, SUBSTR_AFTER (SUBPATH, '.') AS DBID FROM
      SYS_DATABASES.M_VOLUMES);
```

Procedure

1. Log on to the SAP HANA server as operating system user `<sid>adm`.
2. Validate that you have the password for the root key backup file:

```
cd /usr/sap/<sid>/<HDBinstance_no>/exe
./hdbnsutil -validateRootKeysBackup <filename> [--password=<passphrase>]
```

Note

If you don't provide the password on the command line, you will be prompted to enter it.

3. In the system database, stop the tenant database to be recovered.

You can do this in the SAP HANA cockpit or by executing the statement `ALTER SYSTEM STOP <database_name>`.

4. Import the backed-up root keys using the `hdbnsutil` program:

```
cd /usr/sap/<sid>/<HDBinstance_no>/exe
./hdbnsutil -recoverRootKeys <filename>.rkb --dbid=<dbid> --
password=<passphrase> --type=ALL
```

→ Recommendation

We recommend that you do not enter the password on the command line. You will be interactively prompted to enter it. In this way, you avoid unintentionally leaving the password in command history and making it visible in process monitoring tools provided by the operating system.

Note

If you have backed-up root keys to different files, for example according to root key type, you need to execute the command several times.

i Note

- `<dbid>` is the tenant database ID.
- The `<type>` option is the root key type and also accepts the values `PERSISTENCE`, `LOG`, `BACKUP`, and `APPLICATION`. The value `ALL` specifies that root keys of all types are imported. If you do not specify any value for `<type>`, all key types are imported.

Results

The instance SSFS is initialized with the imported root keys. Root keys of the imported type already in the SSFS are overwritten.

Next Steps

Recover the database. For more information, see the section on database recovery.

Related Information

[Stop a Tenant Database \[page 214\]](#)

[Recover a Database \[page 1347\]](#)

7.4.1.11 Use FIPS 140-2 Certified Cryptographic Kernel in CommonCryptoLib

The SAP Cryptographic Library, CommonCryptoLib, supports a FIPS 140-2 compliant cryptographic kernel module, which must be enabled if required.

Prerequisites

You are using CommonCryptoLib patch level 8.4.37 or higher. You can check your version with the following statement: `SELECT * FROM "SYS"."M_HOST_INFORMATION" WHERE KEY LIKE 'crypt%';`

i Note

This statement also shows current version information of your FIPS-compliant crypto kernel if already enabled. If FIPS mode is disabled, the version number is `none`.

Procedure

1. In the database, set the value of the parameter `[cryptography] ccl_fips_enabled` in the `global.ini` configuration file to **true**.
2. Restart the database.

Results

The FIPS 140-2 certified crypto kernel, `libslcryptokernel`, is used instead of the built-in crypto kernel, `libsapcrypto.so`.

If `libslcryptokernel` is not a FIPS 140-2 certified one, the initialization of the library will fail. This means that SAP HANA server processes will not start because of dependent errors in other security functions, for example license errors, SSL errors, and so on.

Related Information

[Modify a System Property in SAP HANA Cockpit \[page 299\]](#)

[Start a Tenant Database \[page 213\]](#)

[Stop a Tenant Database \[page 214\]](#)

[SAP Note 2093286](#)

[SAP Note 2117112](#)

7.4.2 SAP HANA Client Secure User Store (hdbuserstore)

The secure user store (`hdbuserstore`) is a tool installed with the SAP HANA client. Use it to store connection information to SAP HANA systems securely on the client so that client applications can connect to SAP HANA without users having to enter this information. It is typically used by scripts connecting to SAP HANA.

The secure user store allows you to store SAP HANA connection information, including user passwords, securely on clients. In this way, client applications can connect to SAP HANA without the user having to enter host name or logon credentials. You can also use the secure store to configure failover support for application servers in a 3-tier scenario (for example, SAP Business Warehouse) by storing a list of all the hosts that the application server can connect to.

i Note

The secure user store can any be used for all supported clients. The SAP HANA studio however does not use the SAP HANA secure user store, but the Eclipse secure storage. For more information, see the Eclipse documentation.

For more information about the secure user store, see the *SAP HANA Security Guide*.

7.4.2.1 Change the Secure User Store Encryption Key

If you are using the current version of the SAP HANA client, there is no need to change the encryption key of the secure user store. However, if you are using an older version of the SAP HANA client, we recommend changing the encryption key after installation of the SAP HANA client.

Procedure

1. Change the encryption key with the command:

```
hdbuserstore CHANGEKEY
```

The `hdbuserstore` program is available after installation of the SAP HANA client in the following directories:

- `/usr/sap/hdbclient` (Linux/Unix)
- `%SystemDrive%\Program Files\sap` (Microsoft Windows)

A new master encryption key is randomly generated and data in the secure store is re-encrypted with the new key.

2. Verify that the key has been changed with the command:

```
hdbuserstore LIST
```

If the key file `SSFS_HDB.KEY` exists, the time stamp of the file indicates when the key was last successfully changed.

7.4.3 Client-Side Data Encryption

With client-side data encryption, columns that contain sensitive data, such as credit card numbers or social security numbers, are encrypted by using an encryption key accessible only by the client. Client-side encryption makes encryption transparent to applications and column data is encrypted and decrypted on the client-driver, allowing the application to read and write data in cleartext form.

Cleartext data and keys are never available to the SAP HANA server, nor are the cleartext values sent over the network between the client and server. As a result, client-side encryption provides a separation between those who own the data (and can view it) and those who manage the data (but should have no access). Client-side encryption delivers built-in protection of sensitive data from database administrators, cloud administrators, and users who do not need access to cleartext data. Client-side encryption uses both symmetric and asymmetric encryption. Sensitive column data is encrypted with a symmetric column encryption key (CEK) that is encrypted using a client key pair (CKP). A CKP consists of a private key and a public key. The public key is stored on the SAP HANA server and in the `hdbkeystore` on the client's local machine. The private key is stored only in the `hdbkeystore` on the client's local machine.

To access the encrypted data, an application must use a client driver that supports client-side encryption and the client must have access to the CEK that encrypts the column. To distribute these CEKs to clients that need them, key pairs are generated by the clients. The key administrator can then grant access to the CEK, and thus

the encrypted data, by decrypting the CEK with their private key, which is stored in their local hdbkeystore, and creating a copy of the CEK. The CEK copy is encrypted with the public key of the CKP of the user who needs access. The CEK copy for the user is stored in the SAP HANA database. When a user requests a CEK from the server, the server sends the copy encrypted with that user's public key and the client-driver decrypts it using the corresponding private key stored locally.

When writing or reading encrypted data to/from the server, the application must use a prepared statement. When an application issues a parameterized query, the client-driver transparently collaborates with the server to encrypt or decrypt the column data using the key information stored in the SAP HANA server and hdbkeystore. To decrypt the data, the client-driver uses the locally stored private key to decrypt the CEK copy of the user and uses the CEK to access the encrypted column data. The driver encrypts bound parameter values using the CEK (obtained as explained above) before passing the data to the SAP HANA server.

Selecting Deterministic Versus Randomized Encryption

The SAP HANA server never operates on cleartext data stored in encrypted columns. However, depending on the encryption type of the column, some queries on encrypted data are supported. There are two types of client-side encryption: non-deterministic (or randomized) encryption and deterministic encryption.

With deterministic encryption, cleartext is always encrypted into the same ciphertext for a given cleartext and CEK, even over separate executions of the encryption algorithm. This consistency permits operations such as equality and inequality comparisons on encrypted data. Deterministic encryption can result in attackers matching cleartext to encrypted values for a column, especially for columns storing low-cardinality data like status flags, boolean values, or major classifications such as gender.

The key administrator must create their key pair in a secure location, so that no other entities have access to their private key. The key administrator must never operate in the same cloud environment that is hosting the SAP HANA server to avoid having the keys or other sensitive data accessible by the server environment.

When a column is encrypted using the non-deterministic algorithm, the same cleartext yields different ciphertexts, even when encrypting the cleartext several times with the same CEK. Due to the randomness of the ciphertext, non-deterministic encryption is stronger, but limits the types of operations that can be performed with the encrypted data. No operations are possible except for basic inserts, updates and fetches.

Choose the encryption algorithm based on the intended use of the data. Use deterministic encryption for columns that are used as search parameters, for example SSN. Use randomized encryption for data that is used for basic inserts, updates, and fetches.

Feature Details and Considerations

Creating and Altering Tables with Encrypted Columns

- Only columns in row and column tables can be altered to encrypt data.
- Tables with encrypted columns must have a primary key, and the primary key cannot be encrypted.
- Column encryption keys must be located in the same schema as the tables whose columns they are encrypting.

- You cannot alter the table type (row to column, or vice versa) of tables with encrypted columns.
- Only encryption-related column alterations are supported for encrypted columns; you can run any alter column operations on unencrypted columns as usual, without affecting encrypted columns.
- You can only alter the encryption status or encryption key for one column at a time.
- You cannot rename an encrypted column.
- Indexes, foreign keys, partition keys, and check constraints are not supported on encrypted columns
- You cannot include an encrypted column in the RESET BY query of the CREATE | ALTER SEQUENCE statements.
- Encrypted columns are not supported in SQLScript, UDFs, and stored procedures.

Querying Tables with Encrypted Columns

- The application must use prepared statements.
- Queries can perform equality look-ups on columns using deterministic encryption, but cannot perform other operations, such as equality joins, greater than/less than, pattern matching using the LIKE operator, or arithmetical operators.
- Only basic inserts, updates, and fetches are possible on columns using random encryption.
For a detailed description of DML limitations, see *DML Limitations With Client-Side Encryption*.

Supported Column Data Types

You can only encrypt the following column data types:

- BOOLEAN
- DATE, TIME, SECONDDATE, TIMESTAMP
- TINYINT, SMALLINT, INTEGER, BIGINT, REAL, DOUBLE
- VARBINARY
- VARCHAR, NVARCHAR

Data Masking

You cannot configure data masking for client-side encrypted columns.

Calculation Views, OLAP Views, and Join Views

You cannot use encrypted columns in calculation views, OLAP views, or join views.

Related Information

[Create a Column Encryption Key \[page 890\]](#)

[Create a Key Administrator \[page 887\]](#)

[Create a Client Key Pair \[page 888\]](#)

[Supported DML With Client-Side Encryption \[page 883\]](#)

[Grant Client-Side Encryption Privileges \[page 886\]](#)

7.4.3.1 Supported DML With Client-Side Encryption

When querying a table with encrypted columns, there are limits on the types of queries that are supported.

Supported Expressions

Expressions referring to encrypted columns are supported if the encrypted column is referred to by direct column reference expression or alias. You cannot use an encrypted column as input parameter for either built-in or user-defined SQL functions. You cannot use implicit or explicit type casting over encrypted columns.

```
<supported_expression> :=  
<deterministically_encrypted_column>  
| <randomly_encrypted_column>  
| <expression_without_encrypted_column>
```

Supported Predicates

Predicate expressions are supported only if they compare a deterministically encrypted column with a user-specified host variable or with a null value (=, <>, IS NULL, and IS NOT NULL). These predicate expressions can be accompanied with other supported expressions in conjunction and/or in disjunction. Any other predicates referring to encrypted columns are not supported.

```
<supported_predicates> :=  
[NOT] <supported_predicate> [{AND | OR} [NOT] <supported_predicate> [...] ]  
<supported_predicate> ::=  
  <predicate_without_encrypted_column>  
  | <expression_without_encrypted_column> [NOT] IN  
  (<supported_subquery_without_encrypted_column_projection>  
  | [NOT] EXISTS (<supported_subquery_without_encrypted_column_projection>  
  | <d_encrypted_col> { =? | <> ? | IS NULL | IS NOT NULL | [NOT] IN (? [, ?]*) }  
  | ? = <d_encrypted_col>  
  | ? [NOT] IN ( <d_encrypted_col_0> [, <d_encrypted_col_list>]* )  
  
<d_encrypted_col> ::= identifier for a deterministically encrypted column
```

<d_encrypted_col_list> is supported only if all columns in the list share the same type, encryption algorithm and column encryption key.

Both IN and EXISTS with <supported_query> are not supported if <supported_query> contains projections on any encrypted columns.

Supported SELECT and Subquery Expressions

In a SELECT statement, encrypted columns can only be used in SELECT, WHERE, and ON clauses with supported expressions or predicates. Set operators involving encrypted columns are not supported except UNION ALL when it unions deterministically encrypted columns with the same column encryption key.

Encrypted columns are generally supported in SELECT statement subqueries (scalar, multi-row, or derived table expression). However, encrypted columns cannot be used in SELECT clauses of scalar and multi-row/column subqueries as this results in an unsupported comparison between an encrypted column and other expressions contained in the parent query.

```
<supported_query> := <supported_select>
| <supported_select> UNION ALL <supported_select>
<supported_select> :=
SELECT <supported_expression> [, <supported_expression>]*
FROM <table_expression>
[WHERE <supported_predicate>]
[GROUP BY <expression_without_encrypted_column>]
[HAVING <supported_predicate_without_encrypted_column>]
[ORDER BY <expression_without_encrypted_column>]
[LIMIT <expression_without_encrypted_column>]
```

Any subquery expressions not mentioned above are not supported.

Supported Table Expressions

The following table expressions are supported with client-side encryption:

```
<supported_table_expression> :=
<table_reference>
| <table_reference> [, <table_reference>]*
| <table_reference> INNER JOIN <table_reference> ON <supported_predicate>
| <table_reference> LEFT OUTER JOIN <table_reference> ON <supported_predicate>
| <table_reference> RIGHT OUTER JOIN <table_reference> ON <supported_predicate>
| <table_reference> FULL OUTER JOIN <table_reference> ON <supported_predicate>
| <supported_select>
```

Any supported SELECT statement can be placed in a FROM clause.

Supported DML Statements

INSERT, DELETE, UPDATE, and UPSERT statements are partially supported with client-side encryption due to client-side encryption always using the client driver for encryption and decryption. All other operations not listed here are not supported, for example, the MERGE INTO statement.

Supported UPDATE Statements

```
<supported_update_statements> :=
<supported_insert>
| <supported_delete>
| <supported_update>
| <supported_upsert>
```

Supported INSERT Statement

INSERT can be classified into two types, according to the data source: INSERT-VALUES and INSERT-SELECT.

```
<supported_insert> :=
<supported_insert_values>
| <supported_insert_select>
```

```

<supported_insert_values> :=
INSERT <table_reference> [( <column_list> )] VALUES
( <supported_value> [, <supported_value> ] )
<supported_value> := ? | <value_without_encrypted_column>
<supported_insert_select> :=
INSERT <table_reference> [( <column_list> )] SELECT
<supported_expression> [, <supported_expression> ] FROM
{ <table_expression> }

```

The pair of columns and values to insert should have the same column encryption key.

You must use parameter values for encrypted columns as the client-driver cannot access and encrypt values from literals and expressions.

Encrypt the target column and the corresponding source column with the same column encryption key to maintain the integrity of encrypted values.

**Supported
DELETE
Statement**

Support for a DELETE statement depends on whether the predicate is supported:

```

<supported_delete> := DELETE FROM <table_reference> [WHERE
<supported_predicate>]

```

**Supported
UPDATE
Statement**

UPDATE has two different forms: UPDATE FROM WHERE and UPDATE WHERE. Similar to the DELETE statement above, they require a supported predicate in order to be used with client-side encryption. To use the SET clause with client-side encryption, you can only specify parameter values.

```

<supported_update> :=
UPDATE <table_reference> SET <supported_set_clause> WHERE
<supported_where_clause>
| UPDATE <table_reference> SET <supported_set_clause> FROM
<table_reference> [, <table_reference> ] * WHERE <supported_predicate>
<supported_set_clause> := <supported_set_term> [,
<supported_set_term> ]
<supported_set_term> := <column_name> = <supported_value>
<supported_value> := ? | <value_without_encrypted_column>

```

**Supported
UPSERT/
REPLACE
Statement**

The UPSERT subquery is not supported. UPSERT-VALUES behaves like INSERT-VALUES.

```

<supported_upsert> :=
UPSERT | REPLACE <table_reference> [( <column_list> )] VALUES
( <supported_value> [, <supported_value> ] )
<supported_value> := ? | <value_without_encrypted_column>

```

7.4.3.2 Grant Client-Side Encryption Privileges

Grant a user the required privileges to manage client key pairs and column encryption keys, create tables that use specific column encryption keys, or export or import column encryption keys.

Prerequisites

You are logged into your database as a user who can grant the following privileges WITH ADMIN OPTION or WITH GRANT OPTION:

- CREATE CLIENTSIDE ENCRYPTION KEYPAIR system privilege
- DROP CLIENTSIDE ENCRYPTION KEYPAIR system privilege
- CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN privilege

For more information about granting client-side encryption privileges, see the GRANT statement in the *SAP HANA SQL and Systems Views Reference*.

Procedure

Grant a user the required privileges by executing the following statements:

Option	Description	Action
Grant the CREATE CLIENTSIDE ENCRYPTION KEYPAIR system privilege	Grants the user the ability to create client key pairs.	Execute the following GRANT statement, with the option WITH ADMIN OPTION clause if you want <code><user-name></code> to be able to grant the privilege to other users: <pre>GRANT CREATE CLIENTSIDE ENCRYPTION KEYPAIR TO <user-name> [WITH ADMIN OPTION];</pre>
Grant the DROP CLIENTSIDE ENCRYPTION KEYPAIR system privilege	Grants the user the ability to drop client key pairs. Users who have the DROP CLIENTSIDE ENCRYPTION KEYPAIR privilege can drop key-pairs belonging to other users.	Execute the following GRANT statement, with the option WITH ADMIN OPTION clause if you want <code><user-name></code> to be able to grant the privilege to other users: <pre>GRANT DROP CLIENTSIDE ENCRYPTION KEYPAIR TO <user-name> [WITH ADMIN OPTION];</pre>

Option	Description	Action
Grant the CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN schema privilege	Grants the user the ability to create, alter, drop column encryption keys, and create a key copy for a client key pair.	Execute the following GRANT statement, with the option WITH GRANT OPTION clause if you want <code><user-name></code> to be able to grant the privilege to other users: <pre>GRANT CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN ON SCHEMA <schema_name> TO <user-name> [WITH GRANT OPTION];</pre>

Results

The user has been granted the specified privileges.

7.4.3.3 Create a Key Administrator

Create a client-side encryption key administrator who has the privileges to manage column encryption keys.

Prerequisites

You are connected to your schema as the system user.

Context

A key administrator is a user with the CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN privilege. When a key administrator creates a new column encryption key and encrypts it with their public key, no other system user can access it, even other database administrators, unless the key administrator explicitly creates a copy of the column encryption key for them. Users can get access to encrypted data through copies of the column encryption key encrypted by their client key pair. If the key administrator creates a copy for a user who also has the CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN privilege, or is able to assign the privilege to themselves, they can also create key copies for other users. By default, the schema owner has the CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN privilege on their own schema, but cannot create a column encryption key if they do not have the CREATE CLIENTSIDE ENCRYPTION KEYPAIR privilege.

Procedure

1. Open a SQL console for your schema and execute the following statements to create two new users: a data administrator and a key administrator.

```
CREATE USER <database-admin-user-name>;  
CREATE USER <key-admin-user-name>;
```

2. Execute the following statement to grant <database-admin-user-name> the required database administration privileges:

```
GRANT CREATE SCHEMA TO <database-admin-user-name>;
```

3. Execute the following statement to grant <key-admin-user-name> the privileges required to create client key pairs.

```
GRANT CREATE CLIENTSIDE ENCRYPTION KEYPAIR TO <key-admin-user-name>;
```

4. The database administrator creates a schema and grants <key-admin-user-name> the necessary privileges to manage column encryption keys:

```
CREATE SCHEMA <new-schema>;  
GRANT CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN ON SCHEMA <new-schema> TO <key-admin-user-name>;
```

5. (Optional) To give the key administrator the ability to grant the CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN privilege to other users, use the WITH GRANT OPTION:

```
GRANT CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN ON SCHEMA <my-schema> TO <key-admin-user-name> WITH GRANT OPTION;
```

Results

You have created an administrative user with the privileges required to create, alter, and drop column encryption keys. The key administrator uses the ALTER statement to create a key copy for a client key pair to grant access to encrypted data.

7.4.3.4 Create a Client Key Pair

Create a client key pair.

Prerequisites

You are connected to your schema as the key administrator or as a user with the CREATE CLIENTSIDE ENCRYPTION KEYPAIR privilege.

Context

Client Key Pairs (CKPs) are generated by the client-driver. CKPs are asymmetric keys used to distribute column encryption keys to clients.

Procedure

Open a SQL console for your schema and execute the following statement to create a new client key pair:

```
CREATE CLIENTSIDE ENCRYPTION KEYPAIR <key-name> ALGORITHM 'RSA-OAEP-2048';
```

Results

The new client key pair is created with the specified key name and the algorithm RSA-OAEP-2048. The client key pair is stored, along with its key name and UUID, in the hdbkeystore. The public part of the client key pair is stored in the SAP HANA system catalog. CKPs are named database level objects. CKPs are not shared between database systems. If a client accesses multiple databases that support client-side encryption, then it needs a unique CKP for each one.

Since the key administrator's private key is stored in the hdbkeystore on a local computer, all subsequent operations must be performed on that computer unless the key administrator exports the key pair from the hdbkeystore and imports them to a new computer or generates a new key pair for the new computer and creates new key copies of CEKs encrypted with the new key pair.

7.4.3.5 Drop a Client Key Pair

Delete a client key pair.

Prerequisites

You must have the DROP CLIENTSIDE ENCRYPTION KEYPAIR privilege to drop other users' key pairs.

Context

Dropping a key pair drops all the column encryption key copies encrypted with the key pair. But if the column encryption key copy being dropped belongs to the key administrator, then the SAP HANA server ensures that at least one other key administrator column encryption key copy exists for that CEK. If there is no other key administrator key copy for that CEK, then the drop of key pair fails.

Procedure

Open a SQL console for your schema and execute the following statement to drop a client key pair:

```
DROP CLIENTSIDE ENCRYPTION KEYPAIR <keypair-name>;
```

Results

The client key pair and every column encryption key assigned to it are dropped.

7.4.3.6 Create a Column Encryption Key

Create a column encryption key to encrypt columns containing sensitive data.

Prerequisites

You are connected to a schema as the key administrator.

Procedure

Open the SQL console for your schema and execute the following statement to create a new column encryption key:

```
CREATE CLIENTSIDE ENCRYPTION COLUMN KEY <key-name> ALGORITHM '<algorithm-name>'  
  ENCRYPTED WITH KEYPAIR '<keypair-name>;'
```

The specified key pair should belong to the key administrator creating the column encryption key so that key copies can be made.

Results

The new column encryption key is created and stored in the SAP HANA system catalog.

7.4.3.7 Drop a Column Encryption Key

Delete a column encryption key.

Prerequisites

You are connected to a schema as the key administrator.

If any column is encrypted using the specified key, then no key copies are dropped.

Procedure

Open a SQL console and execute the following statement to drop a new column encryption key:

```
DROP CLIENTSIDE ENCRYPTION COLUMN KEY <key-name>;
```

Results

The specified column encryption key and all of its copies are dropped. If the column encryption key copy being dropped belongs to the key administrator, then the SAP HANA server ensures that at least one other key administrator column encryption key copy exists for that CEK. If there is no other key administrator key copy for that CEK, then the drop of the key pair fails.

7.4.3.8 Manage Access to Column Encryption Keys

Grant or revoke access to a column encryption key.

Prerequisites

You are connected to a schema as the key administrator.

If you are granting access, then the user to whom access is being granted has created a key pair.

Context

When a key administrator creates a new column encryption key and encrypts it with their public key, no other user can access it, unless the key administrator explicitly creates a copy of the column encryption key for them.

Procedure

Choose one of the following options:

Option	Action
Grant access to a column encryption key	<p>Execute the following statement to grant the owner of the specified client key pair access to the specified column encryption key:</p> <pre>ALTER CLIENTSIDE ENCRYPTION COLUMN KEY <key-name> ADD KEYCOPY ENCRYPTED WITH KEYPAIR <keypair-name>;</pre>
Revoke access to a column encryption key	<p>Execute the following statement to remove access for the specified client key pair to the specified column encryption key :</p> <pre>ALTER CLIENTSIDE ENCRYPTION COLUMN KEY <key-name> DROP KEYCOPY ENCRYPTED WITH KEYPAIR <keypair-name>;</pre>

Results

Access to the specified column encryption key has been either granted to or revoked from the user with the specified key pair. The server adds the encrypted CEK copy to the SAP HANA catalog when access is granted. When access is removed, the CEK copy is removed from the catalog. If the CEK copy being dropped is a key administrator copy, then the drop fails if at least one other key administrator key copy for that CEK does not exist.

7.4.3.9 Create an Empty Column Encryption Key

Create an empty column encryption key that can be used for schema creation without requiring the creation of actual encryption keys.

Prerequisites

You are connected to a schema as a user with either the CREATE ANY privilege or the CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN privilege.

Context

A CREATE/ALTER TABLE statement can reference an empty column encryption key, but attempts to insert rows into the table fail.

Procedure

Open a SQL console and execute the following statement to create an empty column encryption key:

```
CREATE CLIENTSIDE ENCRYPTION COLUMN KEY <key-name> ALGORITHM '<algorithm-name>'  
HEADER ONLY;
```

Results

An empty column encryption key is created. Only the properties of the column encryption key (CEK) are created along with the UUID. The CEK is not encrypted with a client key pair. Once the headers are replaced with full column encryption key definitions, the SAP HANA server catalog is updated with the encrypted value of the CEK while the CEK UUID remains the same.

7.4.3.10 Populate an Empty Column Encryption Key

A column encryption key is created and the empty column encryption key is updated.

Prerequisites

You are connected to a schema as the key administrator.

Procedure

Execute the following statement to populate an empty column encryption key:

```
ALTER CLIENTSIDE ENCRYPTION COLUMN KEY <key-name>  
  ENCRYPTED WITH KEYPAIR <keypair-name>;
```

The key pair should belong to the key administrator populating the column encryption key so that key copies can be made. If the specified key pair does not belong to the key administrator or does not exist, then the operation fails.

Results

The SAP HANA system catalog is updated with the column encryption key information.

7.4.3.11 Create a Table with an Encrypted Column

Create a table that encrypts a column using client-side encryption.

Prerequisites

You are a user who has been granted the USAGE privilege on the specified column encryption key.

Context

Both the table and the column encryption key must exist in the same schema.

Procedure

Open a SQL console for the schema where the table is being created and execute a statement similar to the following to encrypt a column:

```
CREATE TABLE MyTable (ID INT, Name NVARCHAR(32)
  CLIENTSIDE ENCRYPTION ON WITH myCEK RANDOM);
```

In the example above the column `Name` is encrypted using the column encryption key `myCEK`.

The type of encryption used is `RANDOM`, meaning that similar clear text is encrypted to different cipher text. `NULL` values are replaced with a special value and encrypted to hide the fact that the column is `NULL`.

Results

The table is created and the specified column is encrypted using the specified column encryption key.

7.4.3.12 Change Column Encryption Status

Change an existing encrypted column to unencrypted and vice versa.

Prerequisites

You are a user who has been assigned key copy of the column encryption key used to encrypt or decrypt the data.

Procedure

Choose one of the following options:

Option	Action
Change an existing column from unencrypted to encrypted	Open a SQL console for the schema where the table exists and execute the following statement: <pre>ALTER TABLE <table-name> ALTER (<column-name> ALTER CLIENTSIDE ENCRYPTION ON WITH <column-encryption- key> [RANDOM DETERMINISTIC]);</pre>

Option	Action
	<p>The type of encryption used is RANDOM, meaning that similar cleartext is encrypted to different ciphertext. NULL values are replaced with a special value and encrypted to hide the fact that the column is NULL.</p> <p>Alternatively, you can use DETERMINISTIC encryption, which means that any cleartext value is encrypted into the same ciphertext value.</p>
<p>Change an existing column from encrypted to unencrypted</p>	<p>Open a SQL console for the schema where the table exists and execute the following statement:</p> <pre data-bbox="501 656 1398 734">ALTER TABLE <table-name> ALTER (<column-name> ALTER CLIENTSIDE ENCRYPTION OFF);</pre>
<p>Add an encrypted column to an existing table</p>	<p>Open a SQL console for the schema where the table exists and execute the following statement:</p> <pre data-bbox="501 853 1398 929">ALTER TABLE <table_name> ADD (<column_name> <data_type> CLIENTSIDE ENCRYPTION ON WITH <key_name> RANDOM);</pre>
<p>Cancel client-side encryption for a table</p>	<p>Open a SQL console for the schema where the table exists and execute the following statement:</p> <pre data-bbox="501 1048 1398 1104">ALTER TABLE <table_name> CANCEL CLIENTSIDE ENCRYPTION;</pre> <p>Cancel client-side encryption if you want to roll back your transaction.</p> <p>The only way to stop an encryption/decryption operation is for an administrator to kill the session or for the session owner to interrupt the SQL statement. If the operation is interrupted, you can cancel it by executing the above statement to roll back your transaction.</p>
<p>Continue client-side encryption for a table</p>	<p>Open a SQL console for the schema where the table exists and execute the following statement:</p> <pre data-bbox="501 1395 1398 1451">ALTER TABLE <table_name> CONTINUE CLIENTSIDE ENCRYPTION;</pre> <p>Continue client-side encryption if it has been interrupted by a ALTER SYSTEM CANCEL statement and you want to restart the operation from where the encryption operation was interrupted without re-encrypting the whole column.</p>

Results

You have encrypted or decrypted the specified column.

Related Information

[Abandon a Long-Running ALTER TABLE Operation \[page 899\]](#)

7.4.3.13 Rotate Column Encryption Key

Change the column encryption key (CEK) for a specified encrypted column.

Prerequisites

You are a key administrator who has key copies of the old CEK and the specified new CEK and has SELECT, UPDATE, and ALTER privileges on the table being altered.

Procedure

1. Use the CLIENTSIDE_ENCRYPTION_COLUMN_KEYS system view to identify the key copies for the old CEK.
2. Create key copies of the new CEK for the client key pairs (CKPs) identified above. Do this for every CKP that has the key copy of the old CEK.

Open a SQL console for the schema where the table exists and execute the following statements to change the column encryption key for a specified column:

```
ALTER CLIENTSIDE ENCRYPTION COLUMN KEY <new-column-encryption-key-name>  
ADD KEYCOPY ENCRYPTED WITH KEYPAIR <keypair-name>;
```

3. Alter the table to re-encrypt the specified column with the new CEK:

```
ALTER TABLE <table-name> ALTER <column-name>  
ALTER CLIENTSIDE ENCRYPTION WITH <new-column-encryption-key> RANDOM;
```

Always specify the type of encryption in the ALTER TABLE statement, even if you want to preserve the previous encryption type. The type of encryption used above is RANDOM.

7.4.3.14 Recovering a Client Key Pair

Implementing the following best practices ensures that you can recover your client key pair if the machine on which it is stored crashes, or if the key administrator is unavailable.

Back Up Your Client Key Pairs to a Secondary Secure Key Store

In case the machine storing your client key pairs crashes, or must be offline for a period of time, ensure that your client key pairs are always accessible by implementing the following best practice:

1. Export your key pair(s) from the secure key store (hdbkeystore).

2. Save the exported key pair(s) in a safe location.
3. Import the key pair(s) into the new secure key store (hdbkeystore) once your machine has recovered.

Ensure Key Administrator Redundancy

Ensure that you have a secondary key administrator in case your primary key administrator is unavailable.

1. The secondary key administrator creates a key pair, exports the key pair from the secure key store (hdbkeystore) and saves the key pair in a safe place.
2. The primary key administrator assigns a "recovery" key copy of each of the column encryption keys they manage to the secondary key administrator.
3. When necessary, ensure that the secondary key administrator is granted CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN privilege either by the primary key administrator, or, if the key administrator is unavailable, by another privileged user.

7.4.3.15 Import and Export Column Encryption Keys

Prerequisites

For exporting a column encryption key (CEK), you must have the following privileges:

- EXPORT system privilege
- USAGE privilege on all CEKs being exported

For importing a CEK and you must have the following privileges:

- IMPORT system privilege
- CREATE CLIENTSIDE ENCRYPTION KEYPAIR system privilege
- CLIENTSIDE ENCRYPTION COLUMN KEY ADMIN schema privilege

Importing encrypted data into an existing table requires the UUIDs of the CEKs of the exported data to match the UUIDs of the CEKs of existing table columns.

Context

Exporting a CEK is equivalent to exporting all CEK copies of the CEK. The export of a CEK copy includes an export of the public key of all the relevant client key pairs.

Procedure

Choose one of the following options:

Option	Action
Export a CEK	Execute the following statement to export a CEK: <pre>EXPORT CLIENTSIDE ENCRYPTION COLUMN KEY <column_encryption_key_name> AS CSV INTO '<filepath>';</pre>
Import a CEK	Execute the following statement to import a CEK: <pre>IMPORT CLIENTSIDE ENCRYPTION COLUMN KEY <column_encryption_key_name> FROM '<filepath>';</pre>

7.4.3.16 Abandon a Long-Running ALTER TABLE Operation

Abandon a long-running operation and then revert the table to its pre-encrypted state, or continue encrypting the table.

Context

All pending transactions on a connection where an ALTER TABLE statement is executed are committed before starting the ALTER TABLE operation.

Procedure

1. An ALTER TABLE operation can fail if it encounters an error. Alternately, you can explicitly abandon a long-running ALTER TABLE operation by choosing one of the following actions:

Option	Action
Explicitly issue a cancel session request by executing a statement similar to the following:	Execute the following statement: <pre>ALTER SYSTEM CANCEL [WORK IN] SESSION <connection_id>;</pre>

Option	Action
Issue a cancel request via the client.	The request depends on the client being used. For example, use <code>cancel ()</code> for JDBC or <code>SQLCancel ()</code> for ODBC.

- Once the operation has been abandoned, choose to either cancel or continue client-side encryption for the table:

Option	Action
Cancel client-side encryption	Execute the following statement: <pre>ALTER TABLE <table_name> CANCEL CLIENTSIDE ENCRYPTION;</pre>
Continue client-side encryption:	Execute the following statement: <pre>ALTER TABLE <table_name> CONTINUE CLIENTSIDE ENCRYPTION;</pre>

Results

You have abandoned a long-running query and either canceled or continued encryption on the specified table.

7.5 Managing Client Certificates

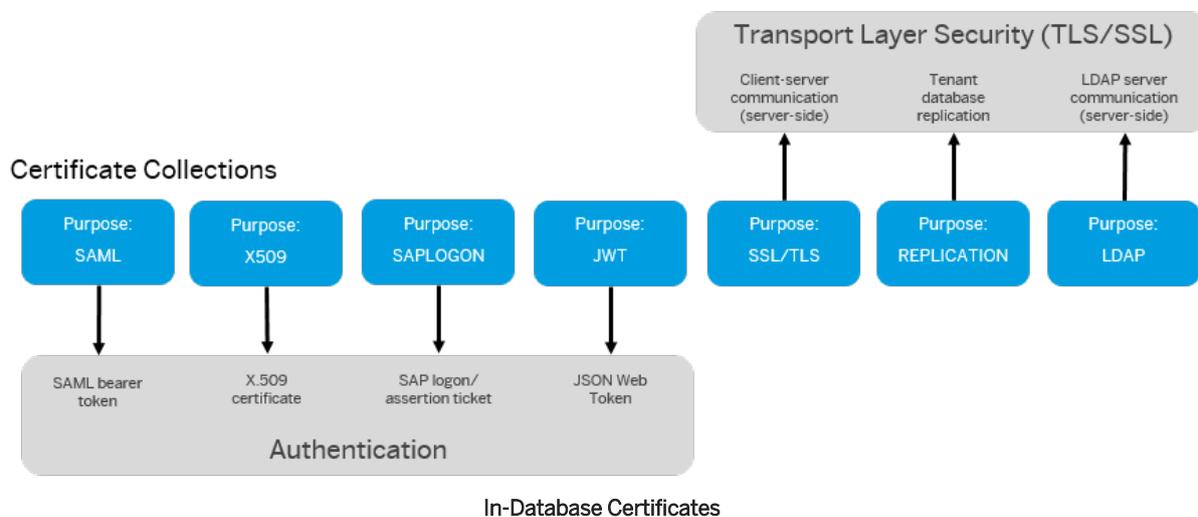
SAP HANA uses X.509 client certificates as the basis for securing internal and external communication channels, as well as for several user authentication mechanisms. Certificates can be stored and managed in files in the file system and in some cases directly in the SAP HANA database.

Certificate Management in the Database

All certificate-based user authentication mechanisms in SAP HANA, as well as secure communication between SAP HANA and clients that access the SQL interface of the database rely on X.509 client certificates for authentication and verifying digital signatures. For ease of management, it's possible to store these certificates and configure their usage directly in the SAP HANA database.

In addition, in-database certificates must be used to secure communication during the process of copying or moving a tenant database between two systems, and to secure communication between SAP HANA and an LDAP server being used for user authentication and authorization.

The following figure shows for which purposes in-database certificates stored in certificate collections can be used. In-database certificates and certificate collections can be fully managed in the SAP HANA cockpit.



Certificate Management in the File System

Although we recommend using in-database storage where possible, you can store and manage certificates in trust and key stores located in the file system, in so-called personal security environments or PSEs.

⚠ Caution

By default, the same PSE in the file system is shared by all databases for all external communication channels (including HTTP) and certificate-based authentication. Different PSEs must be explicitly configured for tenant databases.

→ Recommendation

You can migrate certificates from file-system based storage to in-database storage. If you do migrate certificates in the file system to the database, delete all related files from the file system to avoid any potential conflicts. For more information, see SAP Note 2175664.

However, not all certificates can be stored in the database, in particular the certificates required to secure internal communication channels using the system public key infrastructure (system PKI), and HTTP client access using SAP Web Dispatcher. These certificates are contained in PSE files located in the file system.

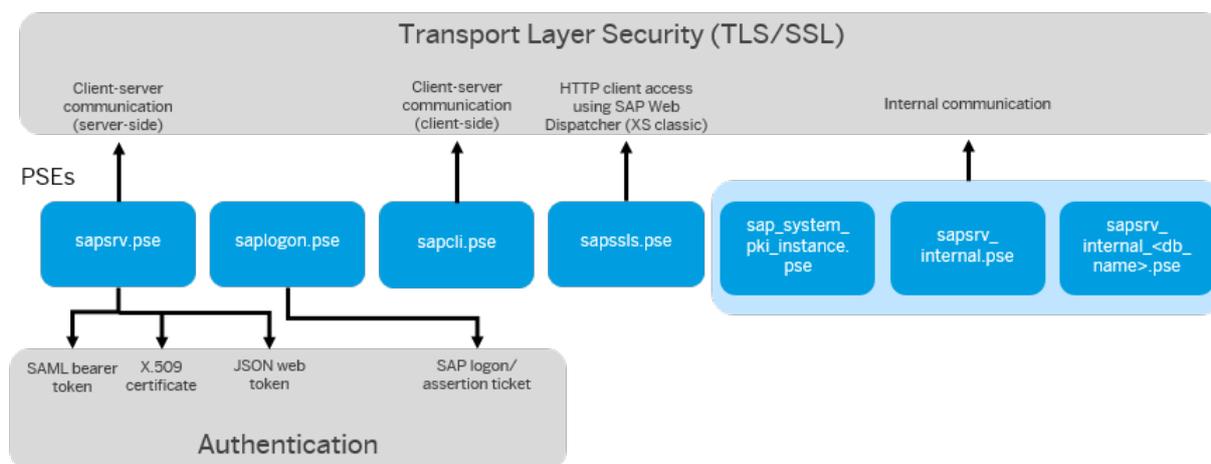
⚠ Caution

Do not delete these files from the file system.

The following figures shows for which purposes certificates stored in PSEs in the file system are possible. These PSEs are available by default and can be managed using for example the SAP Web Dispatcher administration tool or the SAPGENPSE tool, both of which are delivered with SAP HANA. If you are using OpenSSL, you can also use the tools provided with OpenSSL.

i Note

OpenSSL is deprecated. If you are using OpenSSL, migrate to CommonCryptoLib. For more information, see SAP Note 2093286.



Default File-Based PSEs

[In-Database Certificate Management Workflow \[page 903\]](#)

Managing certificates in the SAP HANA database follows a typical workflow. A full separation of duties is possible through user authorization. The full workflow is supported by the SAP HANA cockpit.

[Client Certificates \[page 904\]](#)

X.509 client certificates required for certificate-based authentication and secure communication between SAP HANA and clients that access the SQL interface of the database can be stored and managed directly in the SAP HANA database.

[Certificate Collections \[page 904\]](#)

A certificate collection (or PSE) is a secure location where the public information (public-key certificates) and private information (private keys) of the SAP HANA server are stored. A certificate collection may also contain the public information (public-key certificates) of trusted communication partners or root certificates from trusted Certification Authorities.

[View Certificates in the Certificate Store \[page 906\]](#)

You can view certificates stored in the database on the *Certificate Store* page of the SAP HANA cockpit.

[View Certificate Collections \[page 907\]](#)

You can view the certificate collections available in the database on the *Certificate Collections* page of the SAP HANA cockpit.

[Import a Trusted Certificate into the Certificate Store \[page 909\]](#)

You can store the public-key certificates of trusted communication partners, as well and the root certificates of trusted Certification Authorities directly in the SAP HANA database. You do this on the *Certificate Store* page of the SAP HANA cockpit.

[Create a Certificate Collection \[page 910\]](#)

You can create a certificate collection on the *Certificate Collections* page. Then, you add the relevant trusted certificates and if necessary, the server certificate.

[Set the Purpose of a Certificate Collection \[page 912\]](#)

You specify the purpose of a collection on the *Certificate Collections* page, for example SAML user authentication. A collection may have only one purpose and a purpose may only be served by one collection.

[Export a Client Certificate \[page 913\]](#)

You can export the contents of a client certificate available in the certificate store. For example, you may need to export the SAP HANA server certificate to set up a trust relationship with trusted clients.

All administration tasks related to in-database certificate management can be performed using SQL.

Related Information

[Copying and Moving Tenant Databases Between Systems \[page 1004\]](#)

[SAP Note 2175664](#)

[SAP Note 2093286](#)

7.5.1 In-Database Certificate Management Workflow

Managing certificates in the SAP HANA database follows a typical workflow. A full separation of duties is possible through user authorization. The full workflow is supported by the SAP HANA cockpit.



In-Database Certificate Management Workflow

1. A user with **CERTIFICATE ADMIN** privilege imports into the certificate store the public-key certificates of trusted communication partners, as well as the root certificates of trusted Certification Authorities.
2. A user with **TRUST ADMIN** privilege:
 1. Creates the required certificate collections.
 2. Adds trusted certificates from the certificate store to certificate collections.
 3. Adds the SAP HANA server certificate(s) to those collections that will be used for server authentication (for example, secure client-server communication over JDBC/ODBC).

3. A user with `USER ADMIN`, `SSL ADMIN`, or `DATABASE ADMIN` privilege sets the purpose of individual collections. Which privilege is required depends on the purpose being set.

7.5.2 Client Certificates

X.509 client certificates required for certificate-based authentication and secure communication between SAP HANA and clients that access the SQL interface of the database can be stored and managed directly in the SAP HANA database.

Certificates stored in the SAP HANA database can be used for:

- Trust validation
Certificates used for trust validation are the public-key certificates of trusted communication partners or root certificates from trusted Certification Authorities. These certificates contain the public part of a user's or component's public and private key pair.
- Server authentication
Certificates used for server authentication are the public-key certificates of the SAP HANA server used to identify the server to connecting clients. In addition to the public-key information of the server, these certificates contain the server's private keys, as well as the intermediate certificates that complete the trust chain from the server certificate to the root certificate that the communication partner (client) trusts.

i Note

Private keys are stored securely using the internal application encryption service of the SAP HANA database. For more information, see *Server-Side Data Encryption* in the *SAP HANA Security Guide*.

Once they have been imported into the database, certificates can be assigned to certificate collections. Certificate collections are also created and managed directly in the database, where they serve a unique purpose (either secure client-server communication or a certificate-based authentication mechanism).

i Note

Although we recommend creating and managing both certificates and certificate collections in the database, files containing certificates may also be stored in the file system.

Related Information

[Certificate Collections \[page 904\]](#)

7.5.3 Certificate Collections

A certificate collection (or PSE) is a secure location where the public information (public-key certificates) and private information (private keys) of the SAP HANA server are stored. A certificate collection may also contain the public information (public-key certificates) of trusted communication partners or root certificates from trusted Certification Authorities.

Certificate collections can be created and managed as database objects directly in the SAP HANA database.

Certificate collections uniquely serve one of the following purposes in the database in which they exist:

i Note

Although we recommend creating and managing both certificates and certificate collections in the database, files containing certificates may also be stored in the file system.

- User authentication based on:
 - SAML assertions
 - X.509 certificates
 - Logon and assertion tickets
 - JSON Web Token (JWT)
- Client-server communication over JDBC/ODBC secured using TLS/SSL
- Database replication for the purposes of copying or moving a tenant database to another system
- Communication between SAP HANA and an LDAP server being used for user authentication and authorization

Only one certificate collection may serve one of these purposes at any given time.

The client certificates required for each purpose are assigned to the corresponding certificate collection from the in-database certificate store. A certificate can be assigned to more than one certificate collection.

Certificates used for server authentication, that is certificates that include the private key of the server, need only be assigned to the certificate collection used for secure client-server communication.

Ownership of Certificate Collections

A certificate collection is a database object created in runtime. It is therefore owned by the database user who creates it. If a certificate collection is in use, in other words it has been assigned one of the above purposes, it is not possible to change it (for example, add or remove certificates) or to delete it. However, if the owner of the certificate collection is deleted, the certificate collection will be deleted **even if it currently in use**.

⚠ Caution

The deletion of a certificate collection that is assigned a purpose could render the database unusable. For example, if TLS/SSL is being enforced for all client connections and the certificate collection used for TLS/SSL is deleted, no new client connections to the database can be opened.

Related Information

[Copying and Moving Tenant Databases Between Systems \[page 1004\]](#)

7.5.4 View Certificates in the Certificate Store

You can view certificates stored in the database on the [Certificate Store](#) page of the SAP HANA cockpit.

Prerequisites

You have the system privilege `CERTIFICATE ADMIN` or `TRUST ADMIN`.

Procedure

On the [Overview](#) page, choose the security quick link [Manage certificates](#).

The [Certificate Store](#) page opens. All certificates in the certificate store are listed. If you want to view the full details of a certificate, simply click it. For more information, see [Certificate Details](#).

If the certificate is used in one or more certificate collections, you can navigate to the [Certificate Collections](#) page by clicking the collection name in the [Used In](#) column.

i Note

You will only see the certificate collection if you have the object privilege `ALTER`, `DROP`, or `REFERENCES` on the collection.

7.5.4.1 Certificate Details

On the [Certificate Store](#) page of the SAP HANA cockpit you can view the details of all certificates in the certificate store of the SAP HANA database.

Field	Description
Issued To (CN)	Common name of the person or entity identified by the certificate
Issued To (DN)	Distinguished name of the person or entity identified by the certificate
Issued By (CN)	Common name of the entity that verified the information and issued the certificate
Issued By (DN)	Distinguished name of the entity that verified the information and issued the certificate
Issued On	Date on which the certificate was issued
Expires On	End of certificate's validity

Field	Description
<i>Used In</i>	The certificate collections to which the certificate has been assigned
<i>Version</i>	X.509 version (as specified in the corresponding RFC)
<i>Public Key Algorithm</i>	Public key algorithm
<i>Public Key Length</i>	Public key length
<i>Signature Algorithm</i>	The cryptographic algorithm used to sign the certificate
<i>Basic Constraints</i>	Whether the certificate belongs to a certification authority (CA)
<i>Fingerprint</i>	The hash of the entire certificate, used as a unique identifier in the certificate store
<i>Serial Number</i>	Serial number assigned by the certificate issuer

Related Information

[Certificate Collection Details \[page 908\]](#)

7.5.5 View Certificate Collections

You can view the certificate collections available in the database on the [Certificate Collections](#) page of the SAP HANA cockpit.

Prerequisites

You have the system privilege `CATALOG READ` and either `TRUST ADMIN`, `USER ADMIN`, or `SSL ADMIN`.

Procedure

On the [Overview](#) page, choose the security quick link [Manage certificate collections](#). The [Certificate Collections](#) page opens. All existing collections are listed on the left. To see more detailed information about a specific collection on the right, simply select it. For more information, see [Certificate Collection Details](#).

i Note

In back-end terminology, certificate collections are referred to as personal security environments (PSEs).

Related Information

[Certificate Collections](#) [page 904]

7.5.5.1 Certificate Collection Details

On the [Certificate Collections](#) page of the SAP HANA cockpit, you can view the details of all certificate collections in the SAP HANA database.

Field	Description
Purpose	<p>Purpose of the collection:</p> <ul style="list-style-type: none">• User authentication based on:<ul style="list-style-type: none">◦ SAML assertions◦ X.509 certificates◦ Logon and assertion tickets◦ JSON Web Token (JWT)• Client-server communication over JDBC/ODBC secured using TLS/SSL• Database replication for the purposes of copying or moving a tenant database to another system• Communication between SAP HANA and an LDAP server being used for user authentication and authorization
Provider	SAML identity provider if the collection purpose is SAML
Private Key	<p>Indicates whether or not a private key has been set for the collection</p> <p>Only a collection with the purpose SSL requires a private key. This is the key that the SAP HANA server uses to identify itself to connecting clients.</p>
Created By	Database user who created the collection
Comment	Optional comment

Field	Description
Certificates	<p>Certificates assigned to the collection</p> <p>The function of each certificate in the certificate collection is indicated. The following functions are possible:</p> <ul style="list-style-type: none"> • TRUST The certificate is the public-key certificate of a trusted communication partner. • PERSONAL The certificate is a server certificate belonging to the SAP HANA system and contains a private key. • CHAIN The certificate is an intermediate certificate that is part of the trust chain from the server certificate to the root certificate that the communication partner (client) trusts. <p>For more information about the other certificate fields, see Certificate Details.</p>

Related Information

[Certificate Details \[page 906\]](#)

7.5.6 Import a Trusted Certificate into the Certificate Store

You can store the public-key certificates of trusted communication partners, as well and the root certificates of trusted Certification Authorities directly in the SAP HANA database. You do this on the [Certificate Store](#) page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege System privilege `CERTIFICATE ADMIN`.
- The certificate that you want to add is available on your client in PEM format.

Procedure

1. On the [Overview](#) page, choose the security quick link [Manage certificates](#).
The [Certificate Store](#) page opens, listing all certificates already in the certificate store.

2. Import the certificate:
 - a. Click *Import*.
 - b. Specify the location of the certificate file on your client or paste the content of the file.
 - c. Click *OK*.

The certificate is imported into the database and appears in the list of certificates in the certificate store. You can see the content of the certificate by navigating to its details view. For more information, see *Certificate Details*.

Results

The certificate is available for assignment to one or more certificate collections.

Related Information

[Certificate Details \[page 906\]](#)

7.5.7 Create a Certificate Collection

You can create a certificate collection on the *Certificate Collections* page. Then, you add the relevant trusted certificates and if necessary, the server certificate.

Prerequisites

- You have the system privilege `TRUST ADMIN`.
- The certificates you want to add to the collection are in the certificate store. For more information, see *Add a Certificate to the Certificate Store*.
- If you plan to add a server certificate to the collection, it is available on your client in PEM format.

Procedure

1. On the *Overview* page, choose the security quick link *Manage certificate collections*.
The *Certificate Collections* page opens. All existing collections are listed on the left.
2. Create a new collection by clicking the **+Add** button in the footer toolbar and entering the name of the collection.
The collection is created and appears in the list of collections on the left.

⚠ Caution

You are the owner of the certificate collection. If your database user is deleted, the collection will also be deleted even if it currently in use. This could render the database unusable, for example, if SSL is being enforced for all client connections.

3. Add a trusted certificate by clicking [Add Certificate](#) and then selecting the certificate.

All certificates in the certificate store are available for selection. You can select more than one.

The trusted certificate is added to the collection. It has the function *TRUST*.

4. Optional: Add the server certificate.

In addition to the public-key certificates of trusted communication partners, you can add the certificate of the SAP HANA server. This certificate contains the server's private key, as well as the intermediate certificates that complete the trust chain from the server certificate to the root certificate that the communication partner (client) trusts. The server certificate is necessary if the collection will be used for a purpose that includes server authentication (for example, purpose *SSL*). To add a server certificate, proceed as follows:

- a. Click [Set Own Certificate](#).
- b. Specify the location of the certificate file on your client or paste the content of the file.
- c. Click *OK*.

As a result:

- The server certificate is added to the collection. It has the function *PERSONAL*.
- Any intermediate certificates that are part of the trust chain from the server certificate to the root certificate are also added. They have the function *CHAIN*.
- The *Private Key* attribute changes from *Absent* to *Present*.

Next Steps

Set the purpose of the collection.

Related Information

[Set the Purpose of a Certificate Collection \[page 912\]](#)

7.5.8 Set the Purpose of a Certificate Collection

You specify the purpose of a collection on the [Certificate Collections](#) page, for example SAML user authentication. A collection may have only one purpose and a purpose may only be served by one collection.

Prerequisites

- You have the `REFERENCES` privilege on the certificate collection.
- You have the necessary system privilege to set the purpose:
 - For a user authentication purpose, you need `USER ADMIN`.
 - For the purpose `SSL/TLS` (secure client-server communication over JDBC/ODBC), you need `SSL ADMIN`.

i Note

In addition, the server certificate containing the server's private key must be part of the collection.

- For the purpose `DATABASE REPLICATION`, you need `DATABASE ADMIN`.
- For the purpose `LDAP`, you need `LDAP ADMIN`.

Procedure

1. On the [Overview](#) page, choose the security quick link [Manage certificate collections](#).
2. Find and select the collection that you want to set the purpose for.
3. Open the collection for editing by clicking the [Edit](#) button in the footer toolbar.
4. In the [General Information](#) area, select the purpose:

Option	Description
DATABASE REPLICATION	Communication between two systems via external SQL connections for the purposes of copying or moving a tenant database
JWT	User authentication based on JSON Web Token (JWT)
LDAP	Communication between the SAP HANA database and an LDAP server being used for user authentication and authorization
SAML	User authentication based on SAML assertions
SAP LOGON	User authentication based on logon and assertion tickets
SSL/TLS	Client-server communication over JDBC/ODBC secured using SSL/TLS
X509	User authentication based on X.509 client certificates

i Note

Only the purposes that have been enabled by the system administrator are visible, and only those you are authorized for are enabled.

5. Save the collection.

Results

The collection starts being used for the selected purpose immediately. If another collection had been assigned the purpose before, it will no longer be used.

7.5.9 Export a Client Certificate

You can export the contents of a client certificate available in the certificate store. For example, you may need to export the SAP HANA server certificate to set up a trust relationship with trusted clients.

Prerequisites

You have the system privilege `CERTIFICATE ADMIN` or `TRUST ADMIN`.

Procedure

1. On the *Overview* page, choose the security quicklink *Manage certificates*.

The *Certificate Store* page opens. All certificates in the certificate store are listed.

i Note

You can also navigate to certificates through the certificate collection to which they are assigned.

2. Find the certificate you want to export and navigate to the detailed view.
3. Click *Show PEM Representation* in the footer.
4. Export the certificate contents using copy and paste.

7.5.10 SQL Statements and Authorization for In-Database Certificate Management (Reference)

All administration tasks related to in-database certificate management can be performed using SQL.

The following table lists the SQL statements for creating and managing certificates and certificate collections in the SAP HANA database, including the required authorization for each task.

i Note

Certificate collections are referred to as personal security environments (PSEs) in back-end terminology.

To...	Execute the Statement...	With the Authorization...
See certificates in the in-database certificate store	<pre>SELECT * FROM CERTIFICATES</pre> <div data-bbox="603 528 991 696"><p>i Note</p><p>You can also view certificates in the Certificate Store app of the SAP HANA cockpit.</p></div>	System privilege CERTIFICATE ADMIN or TRUST ADMIN If you have object privilege ALTER on a certificate collection, you'll also be able to see the certificates used in this collection.
See certificate collections	<pre>SELECT * FROM PSES</pre> <div data-bbox="603 770 991 938"><p>i Note</p><p>You can also view certificate collections in the Certificate Store app of the SAP HANA cockpit.</p></div>	System privilege TRUST ADMIN If you have object privilege ALTER, DROP, or REFERENCES on a certificate collection, you'll also be able to see this collection.
See which certificates are used in a certificate collection	<pre>SELECT * FROM PSE_CERTIFICATES</pre> <div data-bbox="603 1046 991 1214"><p>i Note</p><p>You can also see this information on the Certificate Store app of the SAP HANA cockpit.</p></div>	Object privilege ALTER, DROP, or REFERENCES on the certificate collection
Add a certificate to the in-database certificate store	<pre>CREATE CERTIFICATE FROM <certificate_content> [COMMENT <comment>]</pre>	System privilege CERTIFICATE ADMIN
Delete a certificate from the in-database certificate	<pre>DROP CERTIFICATE <certificate_id></pre> <div data-bbox="199 1435 587 1603"><p>i Note</p><p>If the certificate has already been added to a certificate collection, it can't be deleted.</p></div>	System privilege CERTIFICATE ADMIN

To...	Execute the Statement...	With the Authorization...
View certificate collections in the database, including the certificates they contain	<pre>SELECT * FROM PSE_CERTIFICATES</pre> <div style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>You can also view certificate collections on the <i>Certificate Collection</i> app of the SAP HANA cockpit.</p> </div>	<p>System privilege CATALOG READ and either TRUST ADMIN, USER ADMIN, or SSL ADMIN</p> <div style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>If you own a certificate collection or you have the object privilege ALTER, DROP, or REFERENCES on a certificate collection, you'll be able to see it without the above privileges.</p> </div>
Create a certificate collection	<pre>CREATE PSE <PSE_name></pre>	System privilege TRUST ADMIN
Add a public-key certificate to a certificate collection	<pre>ALTER PSE <PSE_name> ADD CERTIFICATE <certificate_id></pre>	<ul style="list-style-type: none"> Nothing if you're the owner of the certificate collection Object privilege ALTER on the certificate collection if you're not the owner
Remove a public-key certificate from a certificate collection	<pre>ALTER PSE <PSE_name> DROP CERTIFICATE <certificate_id></pre> <div style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>If the purpose of the certificate collection already been set, then system privilege USER ADMIN or SSL ADMIN is additionally required depending on whether the purpose is user authentication or secure communication.</p> </div>	
Add a private key to a certificate collection	<pre>ALTER PSE <PSE_name> SET OWN CERTIFICATE <certificate_content></pre>	<ul style="list-style-type: none"> Nothing if you're the owner of the certificate collection Object privilege ALTER on the certificate collection if you're not the owner
Set the purpose of a certificate collection	<pre>SET PSE <PSE_name> PURPOSE <PSE_purpose></pre> <div style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>If the purpose of the PSE is SSL/TLS, then it must already have a private key added.</p> </div>	<ul style="list-style-type: none"> USER ADMIN if the purpose is user authentication (SAML, X.509, JWT, or logon tickets) SSL ADMIN if the purpose is secure client-server communication (SSL/TLS) DATABASE ADMIN if the purpose is copying or moving a tenant database between systems LDAP ADMIN if the purpose is LDAP-based user authentication and authorization
	<p>The following PSE purposes are possible:</p> <ul style="list-style-type: none"> DATABASE REPLICATION JWT LDAP SAML SAP LOGON SSL/TLS X509 	

To...	Execute the Statement...	With the Authorization...
Unset the purpose of a certificate collection	<pre>UNSET PSE <PSE_name> PURPOSE <PSE_purpose></pre>	<p>i Note</p> <p>Object privilege REFERENCES on the certificate collection is additionally required if you are not the owner of the collection.</p>
Delete a certificate collection	<pre>DROP PSE <PSE_name></pre>	<ul style="list-style-type: none"> • Nothing, if you're the owner of the certificate collection • Object privilege DROP on the certificate collection, if you're not the owner
	<p>i Note</p> <p>If the certificate collection has already been assigned a purpose, it can't be deleted.</p>	

7.6 Data Anonymization

To enable analytics on data while still keeping the privacy of individuals, data anonymization capabilities are integrated into SAP HANA calculation views. A list of all calculation views that have one or more anonymization node views configured is available in the SAP HANA cockpit.

i Note

SAP HANA provides only features and tools that help customers to implement data protection requirements and facilitate the required discussions between data scientists and data protection officers.

For more information about data anonymization in SAP HANA, see the *SAP HANA Security Guide*.

For more information about modeling calculation views with anonymization node views, see the *SAP HANA Modeling Guide for XS Advanced Model*.

Related Information

[Show Anonymization Views \[page 917\]](#)

7.6.1 Show Anonymization Views

As a data protection officer or data controller, you can see a list of all calculation views in the database with anonymization node views configured.

Prerequisites

You have system privilege CATALOG READ.

Procedure

On the [Overview](#) page, click the [View anonymization report](#) link.

The [Anonymization View](#) page opens. Here, you can see all calculation views in your database that have one or more anonymization node views configured. The following information is available:

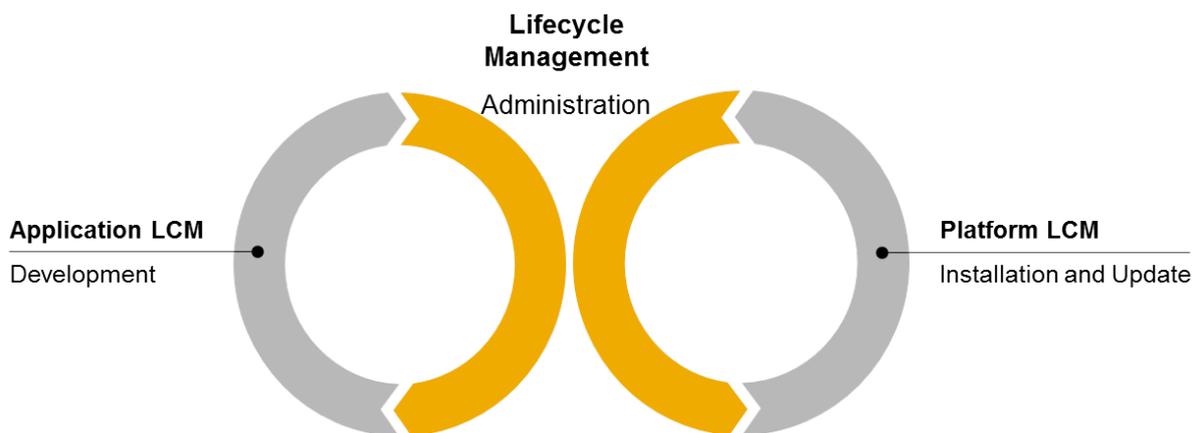
- The name of the calculation view
- The name of the anonymization nodes
- The anonymization method used: k-anonymity or differential privacy
- Configuration values of the relevant method
- Columns in the view, including anonymization information

Next Steps

For documentation purposes, you can download the available information about configured anonymization.

8 SAP HANA Lifecycle Management

SAP HANA lifecycle management covers two aspects: platform lifecycle management for customizing and updating your SAP HANA platform and application lifecycle management for managing SAP HANA content products and transports.



Platform Lifecycle Management Aspects

You can customize platform lifecycle management aspects of your SAP HANA system by accessing the SAP HANA database lifecycle manager from three user interfaces: the graphical user interface, the command-line interface, or the Web user interface in a stand-alone Web browser, in the SAP HANA studio, or via the SAP HANA cockpit.

SAP HANA platform lifecycle management encompasses the installation and update of an SAP HANA server, mandatory components, and additional components, as well as the post-installation configuration. The concepts and procedures for SAP HANA platform installation and update are described in the *SAP HANA Server Installation and Update Guide* on SAP Help Portal.

A number of system configuration features are integrated into the SAP HANA database lifecycle manager, such as:

- The initial configuration of your SAP HANA platform to integrate it into your landscape. For example, by registering it in a system landscape directory, or configuring the inter-service communication.
- Adapting the topology of your SAP HANA platform by adding or removing additional SAP HANA hosts.
- Reconfiguring the system.

System configuration as it pertains to SAP HANA lifecycle management is described in the *SAP HANA Platform Lifecycle Management* section of this *SAP HANA Administration Guide*.

Application Lifecycle Management Aspects

SAP HANA application lifecycle management aspects can be accessed in different user interfaces: an interface that runs as an SAP HANA XS application in a web browser, a command-line tool hdbalm, integrated in SAP HANA studio, or via the SAP HANA cockpit.

SAP HANA application lifecycle management supports you in all phases of the lifecycle of an SAP HANA application or add-on product, from modelling your product structure, through application development, transport, assembly, to installing and updating products that you have downloaded from SAP Support Portal or which you have assembled yourself.

All application lifecycle management tasks are documented in the guide *SAP HANA Application Lifecycle Management* on SAP Help Portal.

System administrators use SAP HANA application lifecycle management mainly to install and update SAP HANA applications or add-on products. Therefore, these tasks are documented in this *SAP HANA Administration Guide*. Tasks related to SAP HANA development are documented in the *SAP HANA Developer Guide - For SAP HANA Studio* (on SAP Help Portal) under *SAP HANA Application Lifecycle Management*.

Related Information

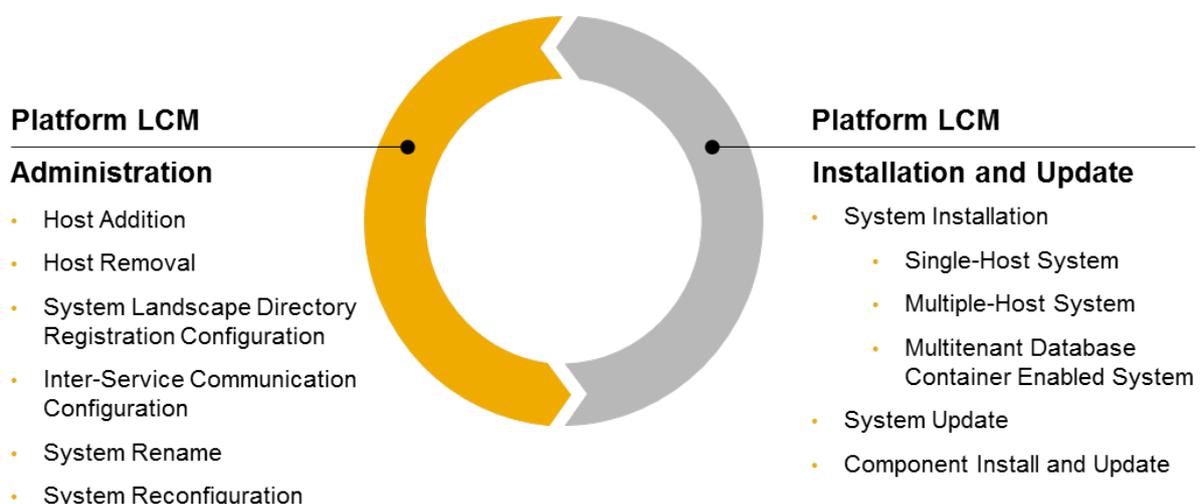
[SAP HANA Platform Lifecycle Management \[page 919\]](#)

[SAP HANA Application Lifecycle Management \[page 951\]](#)

8.1 SAP HANA Platform Lifecycle Management

After the SAP HANA system is installed, it can be configured on the system level.

The SAP HANA platform lifecycle management (LCM) information in this *SAP HANA Administration Guide* details platform administration and configuration. For information about installing and updating the SAP HANA system, see the *SAP HANA Server Installation and Update Guide* on SAP Help Portal.



The SAP HANA database lifecycle manager provides flexibility to accommodate all types of administrators. Before performing administration tasks using the SAP HANA database lifecycle manager, consider reviewing the topic *Using the SAP HANA Platform LCM Tools* to understand the available user interfaces, interaction modes, and parameter entry methods.

You can use the SAP HANA database lifecycle manager to perform the following administration tasks:

- Configure the system
 - Configure a multiple-host system
 - Add one or more hosts to a system
 - Remove one or more hosts from a system
 - Configure a connection to the System Landscape Directory (SLD)
 - Configure inter-service communication
- Change the existing system
 - Change the system identifiers
 - Rename system hosts
 - Change the SID
 - Change the instance number
 - Reconfigure the system
 - Relocate the system to new hardware
 - Copy or clone the system
 - Convert the system to a multitenant database container enabled system

Related Information

[Using the SAP HANA Platform LCM Tools \[page 921\]](#)

[Configuring an SAP HANA System to Connect to the System Landscape Directory \(SLD\) \[page 146\]](#)

[Configuring SAP HANA Inter-Service Communication \[page 1440\]](#)

[Rename an SAP HANA System Host \[page 1034\]](#)

[Change the SID of an SAP HANA System \[page 1036\]](#)

[Change the Instance Number of an SAP HANA System \[page 1039\]](#)

[Relocate the SAP HANA System \[page 999\]](#)

[Copy or Clone an SAP HANA System \[page 1002\]](#)

[Converting an SAP HANA System to Support Tenant Databases \[page 190\]](#)

8.1.1 About the SAP HANA Database Lifecycle Manager (HDBLCM)

The SAP HANA database lifecycle manager (HDBLCM) is used to install, update, or configure an SAP HANA system. You can use the SAP HANA database lifecycle manager in graphical user, command-line, or Web user interface.

8.1.1.1 Using the SAP HANA Platform LCM Tools

The SAP HANA database lifecycle manager (HDBLCM) is used to perform SAP HANA platform lifecycle management (LCM) tasks, including installing, updating, and configuring an SAP HANA system. The SAP HANA database lifecycle manager is designed to accommodate hardware partners and administrators, and so it offers a variety of usage techniques.

The SAP HANA database lifecycle manager is used by means of program interface type, program interaction mode, and parameter entry mode. Before using the SAP HANA database lifecycle manager, you should choose which user interface you prefer to use and how you want to modify the platform LCM task to achieve your desired result. You modify the actions of the platform LCM tools using parameters. Parameters can be modified in a number of ways, for example, in the entry field of a graphical interface, as a call option with the program call, or in a configuration file. These options can be mixed and matched depending on the parameters you need to use and the program interaction mode you choose.

Platform LCM Tools and Program Interaction Modes

	Interactive Mode	Advanced Interactive Mode	Batch Mode
Graphical User Interface	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Command-Line Interface	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Web User Interface	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

The first choice to make is which SAP HANA database lifecycle manager (HDBLCM) interface type you prefer to use. The SAP HANA HDBLCM program can be run as a graphical user interface, a command-line interface, or as Web user interface in a Web browser (the Web user interface is not available for all platform LCM tasks).

Once you've chosen the graphical user, command-line, or Web user interface, you can decide if you prefer to interactively enter parameter values, or give all required parameters with the call to the platform LCM tool, and let it run unattended to completion. Interactive mode is available for all user interfaces, and is the default mode for program interaction. To use interactive mode, you simply call the SAP HANA HDBLCM user interface, and enter parameter values as they are requested by the program. Advanced interactive mode involves entering some parameter values interactively and providing some parameter values as call options or in a configuration file. This is the recommended interaction mode if you'd like to modify parameter default values which are not requested in interactive mode. Batch mode is an advanced platform LCM interaction method because all required parameters must be provided with the call to the LCM program on the command line. Batch mode is designed for large-scale platform LCM tasks, which would be time consuming to perform interactively.

Platform LCM parameters can be entered interactively (only available for interactive mode or advanced interactive mode), as a call option on the command line, or via a configuration file. If you are performing platform LCM tasks in advanced interactive mode, you can choose any of the three parameter entry methods

(or use more than one). If you are using batch mode, you must enter parameter values either as call options to the SAP HANA database lifecycle manager or from a configuration file. The syntax for the parameters as call options can be found in the *Parameter Reference*. The configuration file is generated as a blank template, then edited, and called as a call option.

Related Information

[Use Interactive Mode to Perform Platform LCM Tasks \[page 930\]](#)

[Use Advanced Interactive Mode to Perform Platform LCM Tasks \[page 931\]](#)

[Use Batch Mode to Perform Platform LCM Tasks \[page 933\]](#)

8.1.1.1.1 Choosing the Correct SAP HANA HDBLCM for Your Task

It is important to distinguish between the version of the SAP HANA database lifecycle manager (HDBLCM) that is available on the installation medium and the version that is unpacked during installation, and subsequently used to perform administration and configuration tasks after the SAP HANA system has been installed.

The SAP HANA database lifecycle manager is available in two varieties - an installation medium version to perform installation and update, and a resident version for update and configuration that is unpacked on the SAP HANA host during installation or update. The SAP HANA resident HDBLCM has been designed to be version-compatible. That means, every time you install or update an SAP HANA system, you can be sure that any subsequent configuration tasks performed with the SAP HANA database lifecycle manager will work as expected because the installation or update tool and the configuration tool are of the same version and have been tested together. The SAP HANA resident HDBLCM is located at `<sapmnt>/<SID>/hdb1cm`.

8.1.1.1.2 Performing LCM Tasks by Program Interface

SAP HANA platform lifecycle management tasks can be performed from a graphical, command-line and Web user interface.

Related Information

[Use the Graphical User Interface to Perform Platform LCM Tasks \[page 923\]](#)

[Use the Command-Line Interface to Perform Platform LCM Tasks \[page 924\]](#)

[Using the Web User Interface \[page 925\]](#)

8.1.1.1.2.1 Use the Graphical User Interface to Perform Platform LCM Tasks

SAP HANA platform lifecycle management tasks can be performed from a graphical interface.

Procedure

1. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	<pre>cd <installation medium>/DATA_UNITS/ HDB_LCM_LINUX_X86_64</pre>
Installation Medium (IBM Power Systems)	<pre>cd <installation medium>/DATA_UNITS/ HDB_LCM_LINUX_PPC64</pre>
SAP HANA resident HDBLCM	<pre>cd <sapmnt>/<SID>/hdblcmm</pre>

In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the two SAP HANA database lifecycle manager types, see Related Information.

2. Start the SAP HANA platform lifecycle management tool:

```
./hdblcmmgui
```

3. Enter parameter values in the requested fields. In addition, you can specify parameter key-value pairs as call options or in the configuration file template.

i Note

If parameter key-value pairs are specified as command-line options, they override the corresponding parameters in the configuration file. Parameters in the configuration file override default settings.

Order of parameter precedence:

Command Line > Configuration File > Default

For more information about program interaction modes and parameter values entry methods, see Related Information.

Related Information

[Choosing the Correct SAP HANA HDBLCM for Your Task \[page 922\]](#)

[Entering Platform LCM Parameters as Call Options from the Command Line \[page 937\]](#)

8.1.1.1.2.2 Use the Command-Line Interface to Perform Platform LCM Tasks

SAP HANA platform lifecycle management tasks can be performed from the command line.

Procedure

1. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	<pre>cd <installation medium>/DATA_UNITS/ HDB_LCM_LINUX_X86_64</pre>
Installation Medium (IBM Power Systems)	<pre>cd <installation medium>/DATA_UNITS/ HDB_LCM_LINUX_PPC64</pre>
SAP HANA resident HDBLCM	<pre>cd <sapmnt>/<SID>/hdblcm</pre>

In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the two SAP HANA database lifecycle manager types, see Related Information.

2. Start the SAP HANA platform lifecycle management tool:

```
./hdblcm
```

3. Enter parameter values in one of the following ways.
 - **Interactive parameter entry** - If you call the SAP HANA platform LCM tool only, the program runs in interactive mode. Parameter default values are suggested in brackets, and can be accepted with *Enter*. Otherwise, enter a non-default value, then select *Enter*.
 - **Command-line parameter entry as call options** - If you enter parameter key-value pairs as call options with the call to the SAP HANA platform LCM tool, the program runs in advanced interactive mode and requests values for any parameter values which you didn't specify in the original input. If you entered the batch mode call option, the program runs to completion without any further requests, unless a mandatory parameter was left out of the original input, in which case, the program fails to perform the platform LCM task.
 - **Configuration file parameter entry** - If you enter parameter key-value pairs in the configuration file template, and enter the configuration file path as a call option with the call to the SAP HANA platform LCM tool, the program runs in advanced interactive mode and requests values for any parameter values which you didn't specify in the original input. If you entered the batch mode call option, the program runs to completion without any further requests, unless a mandatory parameter was left out of the original input, in which case, the program fails to perform the platform LCM task.

Note

If parameter key-value pairs are specified as command-line options, they override the corresponding parameters in the configuration file. Parameters in the configuration file override default settings.

Order of parameter precedence:

Command Line > Configuration File > Default

For more information about program interaction modes and parameter values entry methods, see Related Information.

Related Information

[Choosing the Correct SAP HANA HDBLCM for Your Task \[page 922\]](#)

[Performing LCM Tasks by Parameter Entry Method \[page 934\]](#)

[Performing LCM Tasks by Program Interaction Mode \[page 930\]](#)

[Entering Platform LCM Parameters as Call Options from the Command Line \[page 937\]](#)

8.1.1.1.2.3 Using the Web User Interface

SAP HANA platform lifecycle management tasks can be performed using the SAP HANA database lifecycle manager (HDBLCM) Web user interface.

8.1.1.1.2.3.1 About the Web User Interface

The SAP HANA database lifecycle manager (HDBLCM) Web user interface is hosted by the SAP Host Agent, which is installed on the SAP HANA host. When installing or updating the SAP HANA system, as part of the SAP HANA resident HDBLCM configuration, the SAP HANA system deploys its artifacts on the SAP Host Agent, thus enabling the Web user interface.

All Web user interface actions are always performed in the context of an already installed and registered SAP HANA system. In order to access the SAP HANA database lifecycle manager Web user interface you need to log on as the system administrator user `<sid>adm`.

The communication between the Web browser and the SAP Host Agent is always done over HTTPS, which requires that the SAP Host Agent has a secure sockets layer (SSL) certificate (PSE) in its security directory. For more information about SSL certificate handling, see Related Information.

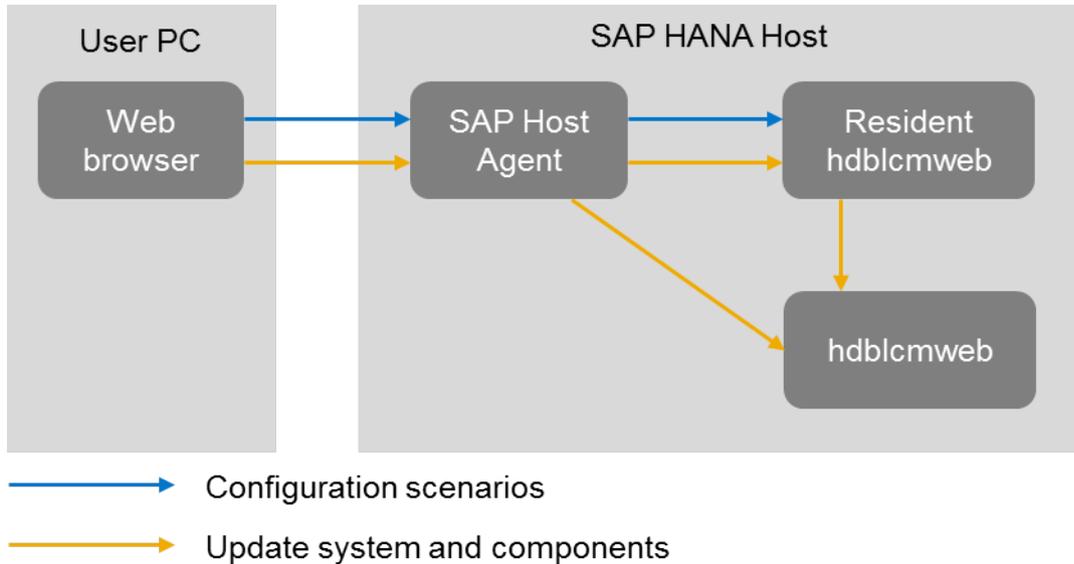
The backend is provided by the special executable `hdb1cmweb`, which is started automatically by the SAP Host Agent as soon as an action is triggered from the Web user interface and terminates after the action completes.

i Note

You should never start `hdb1cmweb` manually. For security reasons, `hdb1cmweb` is always started with system administrator user `<sid>adm` privileges. If you require logging with individual users (to ensure personalized logging), use the SAP HANA database lifecycle manager graphical user or command-line interface.

i Note

Make sure that the system administrator user `<sid>adm` has permissions to read the paths, passed as parameters in the Web user interface (for example, the SAP HANA database installation kit or locations with SAP HANA components).



One platform LCM task, which is worth special attention is the update of the SAP HANA system and components. The SAP HANA system updates are always performed by the installation kit SAP HANA database lifecycle manager in the graphical user and command-line interfaces, (and not the SAP HANA resident HDBLCM). This is because the SAP HANA database lifecycle manager, in the graphical user and command-line interfaces, is not forward compatible. Meaning that only the new version of the tool knows how to update an older system.

On the other hand, all scenarios in the Web user interface are handled by the SAP HANA resident HDBLCM, which is part of the system. For this reason, as a first step before even starting the update process, you are required to enter a location of an SAP HANA database installation kit. After detecting the kit, the update Web user interface is loaded from the installation kit and the installation kit SAP HANA database lifecycle manager starts serving as backend until the update process finishes. It is as if you start the SAP HANA database lifecycle manager directly from the installation kit in graphical user or command-line interface.

Related Information

[Secure Sockets Layer \(SSL\) Certificate Handling \[page 940\]](#)

8.1.1.1.2.3.2 Use the Web User Interface to Perform Platform LCM Tasks

The SAP HANA database lifecycle manager (HDBLCM) can be accessed as a Web user interface in either a standalone browser or in the SAP HANA cockpit.

Prerequisites

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
 - Microsoft Windows
 - Internet Explorer - Version 9 or higher
If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing **Tools > Compatibility View Settings**.
 - Microsoft Edge
 - Mozilla Firefox - Latest version and Extended Support Release
 - Google Chrome - Latest version
 - SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
 - Mac OS - Safari 5.1 or higher

i Note

For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

Context

The Web user interface supports only the following SAP HANA platform lifecycle management tasks:

- View system information
- Update system and components
- Install or update additional components
- Configure System Landscape Directory (SLD) registration
- Configure inter-service communication

When performing installation and update tasks, various parameters can be set in the [Advanced Parameters Configuration](#) dialog. To access the [Advanced Parameters Configuration](#) dialog, click on the gear icon in the footer bar of the SAP HANA HDBLCM Web user interface.

Procedure

Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: <code>https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html</code>
	i Note The URL is case sensitive. Make sure you enter upper and lower case letters correctly.
SAP HANA cockpit	<ol style="list-style-type: none">1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. <code>https://<host_FQDN>:<port></code> i Note FQDN = fully qualified domain name
	<ol style="list-style-type: none">2. Drill down on the name of the system from My Resources or from a group.3. The links in Platform Lifecycle Management each launch additional functionality, giving you expanded capabilities for managing the resource.

Results

The SAP HANA database lifecycle manager is displayed as a Web user interface in either a standalone browser or in the SAP HANA cockpit.

Related Information

[SAPUI5 Developer Guide](#)

[Add an SAP HANA System \[page 122\]](#)

8.1.1.1.2.3.3 Log Off From an SAP HANA System

In the SAP HANA database lifecycle manager (HDBLCM) Web user interface, you can log off from an SAP HANA system and close all connections to the system. To be able to connect to system again, you must log on.

Procedure

- To log off from a system click the [Log out](#) button.
All open connections to the system are closed.

i Note

Currently, this feature is not available for browsers on mobile devices.

8.1.1.1.2.3.4 Troubleshooting the Web User Interface

If you have problems with the Web user interface, see SAP Note 2078425 for steps you can take to troubleshoot and resolve them.

i Note

The Web browser used to render the platform lifecycle management Web user interface in the SAP HANA studio **cannot** be changed via ► [Windows](#) ► [Preferences](#) ► [General](#) ► [Web Browser](#) ►.

Related Information

[SAP Note 2078425](#) 

8.1.1.1.3 Performing LCM Tasks by Program Interaction Mode

SAP HANA platform lifecycle management tasks can be performed in interactive mode, advanced interactive mode and batch mode.

8.1.1.1.3.1 Use Interactive Mode to Perform Platform LCM Tasks

Interactive mode is a method for running SAP HANA platform lifecycle management (LCM) tools which starts the program and requires you to enter parameter values successively before the program is run. Interactive mode is the default mode for the SAP HANA platform LCM tools.

Context

In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the different SAP HANA database lifecycle manager types, see Related Information.

The SAP HANA platform LCM tools offer a wide variety of parameters which can modify the platform LCM task you are performing. Some parameters can be modified in interactive mode when the graphical user, command-line, or Web user interface requests a value for a given parameter.

Procedure

1. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	<pre>cd <installation medium>/DATA_UNITS/ HDB_LCM_LINUX_X86_64</pre>
Installation Medium (IBM Power Systems)	<pre>cd <installation medium>/DATA_UNITS/ HDB_LCM_LINUX_PPC64</pre>
SAP HANA resident HDBLCM	<pre>cd <sapmnt>/<SID>/hdb1cm</pre>

To access the SAP HANA database lifecycle manager Web user interface, see Related Information.

2. Start the SAP HANA platform lifecycle management tool:

Option	Description
Graphical Interface	<code>./hdblcgui</code>
Command-line Interface	<code>./hdblc</code>

To start the SAP HANA platform LCM tools in interactive mode, simply **do not** enter the parameter for batch mode (`--batch` or `-b`) as a call option. You can enter any other required parameters as call options or load a configuration file. The program runs in interactive mode and requests any missing parameters values, which must be verified or changed. You are provided with a summary of parameter values, which you can accept to run the program to completion, or reject to exit the program.

Related Information

[Choosing the Correct SAP HANA HDBLCM for Your Task \[page 922\]](#)

[Use the Web User Interface to Perform Platform LCM Tasks \[page 927\]](#)

8.1.1.1.3.2 Use Advanced Interactive Mode to Perform Platform LCM Tasks

Interactive mode is a method for running SAP HANA platform lifecycle management (LCM) tools which starts the program and requires you to enter parameter values successively before the program is run. If you would like to enter call options not available in interactive mode, or make use of the configuration file, you can use a combination of interactive mode and advanced parameter entry methods.

Context

In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the different SAP HANA database lifecycle manager types, see Related Information.

The SAP HANA platform LCM tools offer a wide variety of parameters which can modify the platform LCM task you are performing. Some parameters can be modified in interactive mode when the graphical user, command-line, or Web user interface requests a value for a given parameter. However, some parameters are not available in interactive mode, and must be specified either as a call option with the call to the platform LCM tool, or from within a configuration file.

Procedure

1. Review which parameters are offered in interactive mode.

If the parameter you want to configure is not available in interactive mode, you have two options. You can either enter the parameter key-value pair as a call option with the call to the platform LCM tool. Alternatively, you can generate a configuration file template, and edit the parameters value in the configuration file. Then call the configuration file as a call option with the call to the platform LCM tool.

Using the configuration file for interactive mode is recommended if you plan to perform the exact same platform LCM task multiple times.

2. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	<pre>cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64</pre>
Installation Medium (IBM Power Systems)	<pre>cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64</pre>
SAP HANA resident HDBLCM	<pre>cd <sapmnt>/<SID>/hdblcm</pre>

3. If you plan to use a configuration file, prepare it with the following steps:

- a. Generate the configuration file template using the SAP HANA platform lifecycle management tool:

Run the SAP HANA platform LCM tool using the parameter `dump_configfile_template` as a call option. Specify an action and a file path for the template. A configuration file template and a password file template are created.

```
./hdblcm --action=<LCM action> --dump_configfile_template=<file path>
```

- b. Edit the configuration file parameters. Save the file.
- c. Edit the password file. Save the file.

4. Start the SAP HANA platform lifecycle management tool:

Start the SAP HANA database lifecycle manager in either the graphical user interface or in the command-line interface, with a call option:

```
./hdblcmgui --<parameter key>=<parameter value>
```

or

```
./hdblcm --<parameter key>=<parameter value>
```

If you are using a configuration file, you must use the call option `--configfile=<file path>`.

Related Information

[Choosing the Correct SAP HANA HDBLCM for Your Task \[page 922\]](#)

8.1.1.1.3.3 Use Batch Mode to Perform Platform LCM Tasks

Batch mode is a method for running the SAP HANA database lifecycle manager which starts the program and runs it to completion without requiring you to interact with it any further. All required parameter values must be passed as call options or from a configuration file.

Prerequisites

- When using batch mode, passwords must either be defined in the configuration file, or passed to the installer using an XML password file and streamed in via standard input. In both cases, it is necessary to prepare the passwords. For more information, see *Specifying Passwords*.

Context

In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the different SAP HANA database lifecycle manager types, see Related Information.

If you are new to performing the desired SAP HANA platform LCM task in batch mode, it is recommended to run some tests before using batch mode in a production environment.

Procedure

1. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	<pre>cd <installation medium>/DATA_UNITS/ HDB_LCM_LINUX_X86_64</pre>
Installation Medium (IBM Power Systems)	<pre>cd <installation medium>/DATA_UNITS/ HDB_LCM_LINUX_PPC64</pre>
SAP HANA resident HDBLCM	<pre>cd <sapmnt>/<SID>/hdblcmm</pre>

2. Start the SAP HANA platform lifecycle management tool:

```
./hdblcmm --batch <additional parameters>
```

or

```
./hdblcmm -b <additional parameters>
```

It is mandatory to provide an SAP HANA system ID (SID) and user passwords during installation. In batch mode, you are restricted to providing these parameter values as call options on the command line (for passwords, by means of an XML file) or in a configuration file. If you don't provide parameter values for the other required parameters, you implicitly accept the default values.

❖ Example

The following example installs the SAP HANA server and client as a single-host system. The SAP system ID and instance number are also specified from the command line. The system passwords are read from a standard input stream by the installer. All other parameter defaults are automatically accepted and no other input is requested in order to complete the installation.

```
cat ~/hdb_passwords.xml | ./hdblcm --batch --action=install --  
components=client,server --sid=DB1 --number=42 --read_password_from_stdin=xml
```

If a configuration file is used in combination with batch mode, an identical system can be installed with a simplified call from the command line. In the following example, passwords are defined in the configuration file, in addition to the action, components, SAP system ID, and instance number.

```
./hdblcm --batch --configfile=/var/tmp/H01_configfile
```

Related Information

[Choosing the Correct SAP HANA HDBLCM for Your Task \[page 922\]](#)

[Use LCM Configuration Files to Enter Parameters \[page 935\]](#)

[Entering Platform LCM Parameters as Call Options from the Command Line \[page 937\]](#)

8.1.1.1.4 Performing LCM Tasks by Parameter Entry Method

SAP HANA platform lifecycle management (LCM) parameter values can be entered in a variety of methods: interactively by iteratively providing values in either the graphical interface of command prompt, as command-line options with the call to the platform LCM tool, or in a configuration file.

SAP HANA platform lifecycle management parameter values allow you to customize your SAP HANA installation, update, or configuration. Parameter values can be entered by **one or more** of the following methods:

- | | |
|--------------------------------|---|
| Interactively (Default) | Using either command line interface, the graphical interface or the Web user interface, most parameters are requested interactively. Default parameter values are proposed in brackets and can be changed or confirmed. Parameters that are not requested (or specified via another method) accept the default value. |
| Command Line Options | Parameters are given in their accepted syntax as a space delimited list after the program call (for example, <code>hdblcm</code> or <code>hdblcmgui</code>). The specified parameters replace the defaults. If any mandatory parameters are excluded, they are requested interactively (unless batch |

mode is specified). All parameters can be entered from the command line. For more details about the accepted parameter syntax, see the inline help output (`--help`) for the individual SAP HANA lifecycle management tool.

Configuration File

The configuration file is a plain text file, for which a template of parameter key-value pairs can be generated, edited, and saved to be called in combination with the program call. If any mandatory parameters are not specified, they are requested interactively (unless batch mode is used). All parameters can be entered in the configuration file. For more information about the configuration file, see Related Information.

i Note

If parameters are specified in the command line, they override the corresponding parameters in the configuration file. Parameters in the configuration file override default settings.

Order of parameter precedence:

Command Line > Configuration File > Default

8.1.1.1.4.1 Entering Platform LCM Parameters Interactively

SAP HANA platform LCM interactive mode is default interaction mode for all platform LCM programs and interfaces.

You can run the graphical, command-line, or Web user interface in interactive mode by simply starting the program, and entering parameter values as they are requested by the program. In interactive mode, parameter default values are suggested in brackets and can be accepted with `Enter`.

Not all parameters are requested in interactive mode. If you would like to configure a parameter not offered in interactive mode, you must enter it as a call option with the call to the platform LCM program, or use corresponding configuration file for the platform LCM task.

8.1.1.1.4.2 Use LCM Configuration Files to Enter Parameters

By defining a prepared configuration file during installation, specified parameter values are used by the SAP HANA platform lifecycle management (LCM) tools to build a customized SAP HANA system.

Context

The configuration file is a plain text file of specified parameters, written in the same syntax as in the command line (except without the leading two dashes `--`). A configuration file template can be generated, edited, and saved to be called with the call to the SAP HANA database lifecycle manager (HDBLCM).

The configuration file template provides a brief, commented-out summary of each parameter. Each parameter is set to its default value.

Procedure

1. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	<pre>cd <installation medium>/DATA_UNITS/ HDB_LCM_LINUX_X86_64</pre>
Installation Medium (IBM Power Systems)	<pre>cd <installation medium>/DATA_UNITS/ HDB_LCM_LINUX_PPC64</pre>
SAP HANA resident HDBLCM	<pre>cd <sapmnt>/<SID>/hdblcm</pre>

In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the two SAP HANA database lifecycle manager types, see Related Information.

2. Generate the configuration file template using the SAP HANA platform lifecycle management tool:

Run the SAP HANA platform LCM tool using the parameter `dump_configfile_template` as a call option. Specify an action and a file path for the template. A configuration file template and a password file template are created.

```
./hdblcm --action=<LCM action> --dump_configfile_template=<file path>
```

3. Edit the configuration file parameters. Save the file.

It is recommended that at least the SAP system ID (`sid`) and the instance number (`number`) are uniquely defined. There are several required parameters, that are provided default values in case they are not customized. For more information, refer to the default values.

Some file path parameters have automatic substitution values as part of the default file path, using the `sid` (SAP HANA system ID) and `sapmnt` (installation path) parameters, so that the substituted values create file paths that are unique and system-specific. For example, the default for the data file path is: `datapath=/hana/data/${sid}`, where `sid` is automatically replaced by the unique SAP HANA system ID.

4. Start the SAP HANA platform lifecycle management tool:

Run the SAP HANA platform LCM tool using the parameter `configfile` as a call option. Specify the file path of the edited template.

```
./hdblcm --configfile=<file path>
```

You can specify the path to a directory in which custom configuration files are saved using the parameter `custom_cfg` as a call option.

Related Information

[Choosing the Correct SAP HANA HDBLCM for Your Task \[page 922\]](#)

8.1.1.1.4.3 Entering Platform LCM Parameters as Call Options from the Command Line

Call options are available for every SAP HANA platform LCM program.

You can use call options for a number of reasons:

- The parameter is not available in interactive mode, but can be entered as a call option.
- You are using batch mode.
- You are using a configuration file, but would like to override a parameter in the configuration file with a new value.
- You are installing an SAP HANA multiple-host system from the command line.

A call option is entered with the following notation:

```
./<program call> --<parameter1 key>=<parameter1 value> --<parameter2 key>=<parameter2 value>
```

Call options start with a double dash (--) if they are written in long-form syntax. Some parameters also have short-form syntax, in which they are preceded with a single dash (-). For more information about call option syntax, see the *Parameter Reference* topics.

8.1.1.1.5 Executing Platform LCM Tasks

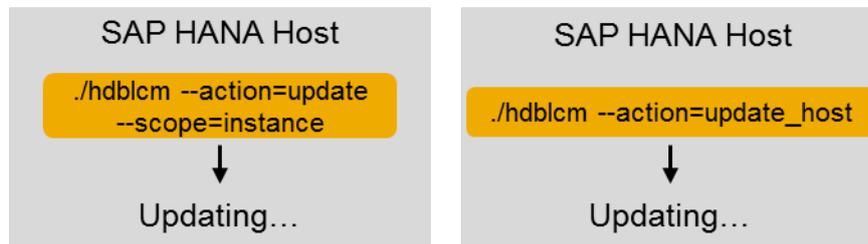
SAP HANA platform lifecycle management tasks can be performed on multiple-host systems centrally, by running the SAP HANA database lifecycle manager (HDBLCM) from any worker host and using remote execution to replicate the call on all remaining system hosts. Otherwise, the platform LCM tasks can be executed first on a worker host, and then re-executed manually on each remaining host. This method is considered decentralized execution.

The following is an example of an SAP HANA system update performed centrally and decentrally.

Centralized Execution



Decentralized Execution



Related Information

[Centralized Execution of Platform LCM Tasks \[page 938\]](#)

[Decentralized Execution of Platform LCM Tasks \[page 943\]](#)

8.1.1.1.5.1 Centralized Execution of Platform LCM Tasks

SAP HANA platform lifecycle management (LCM) tasks can be performed centrally on multiple-host SAP HANA systems in a number of ways depending on the available certificate keys and the remote execution configuration.

8.1.1.1.5.1.1 Using Secure Shell (SSH) to Execute Platform LCM Tasks

An SAP HANA system must be installed with root user credentials. During installation a secure shell (SSH) key is configured so that future platform LCM tasks can be performed remotely on multiple-host SAP HANA systems without requiring the root user password.

By default, the SAP HANA database lifecycle manager (HDBLCM) uses SSH during SAP HANA system installation or update. In order to use SSH, the SFTP subsystem must be active. Install the SAP Host Agent on

all system hosts to perform platform LCM tasks without root credentials. Once the SAP Host Agent is installed, it is used to perform any platform LCM tasks executed from the Web user interface or as the system administrator user `<sid>adm`.

i Note

Platform LCM tasks cannot be executed remotely via SSH as the system administrator user `<sid>adm`.

Related Information

[SAP Note 1944799](#)

[SAP Note 2009879](#)

8.1.1.1.5.1.2 Using SAP Host Agent to Execute Platform LCM Tasks

Platform LCM tasks can be executed without root credentials by using the SAP Host Agent. The SAP Host Agent is installed and updated by default during SAP HANA system installation and update.

The SAP HANA database lifecycle manager (HDBLCM) relies on the SAP Host Agent for the following functionality to work:

- Execution as the system administrator user `<sid>adm`
- Connectivity to remote hosts via HTTPS (when no SSH or root user credentials are available)
- Execution from the SAP HANA database lifecycle manager Web user interface

i Note

The SAP HANA cockpit uses the SAP Host Agent to execute tasks as the system administrator user `<sid>adm`, for example, stopping and starting the system, or troubleshooting a system experiencing performance problems.

If execution on the remote hosts is done via SSH (default, `--remote_execution=ssh`), the SAP HANA database lifecycle manager is able to connect to a remote host via SSH and install and configure the SAP Host Agent. In contrast, the remote execution via SAP Host Agent (`--remote_execution=saphostagent`) requires that the SAP Host Agent is installed and configured on all involved hosts in advance, which includes:

- Install SAP Host Agent
- Configure a Secure Sockets Layer (SSL) certificate for the SAP Host Agent, so that the HTTPS port 1129 is accessible. For more information about SSL configuration for the SAP Host Agent, see Related Information. If you don't want to configure HTTPS, it is also possible to use the call option `--use_http`. It tells the SAP HANA database lifecycle manager to communicate with the SAP Host Agent via HTTP. During the addition of new hosts to an SAP HANA system (also during the installation of a multiple-host system), the HTTPS of the SAP Host Agent is automatically configured by the SAP HANA database lifecycle manager.

⚠ Caution

Use the call option `--use_http` with caution, because passwords are also transferred in plain text via HTTP.

Related Information

[SSL Configuration for the SAP Host Agent](#)

8.1.1.1.5.1.2.1 Secure Sockets Layer (SSL) Certificate Handling

To enable secure communication with the SAP Host Agent over HTTPS, the SAP Host Agent needs a secure sockets layer (SSL) certificate in its security directory. This certificate is also used by the SAP HANA database lifecycle manager (HDBLCM) Web-based user interface because the Web pages are served by the SAP Host Agent.

The SAP HANA database lifecycle manager handles certificate management during system installation, update, or rename, as well as during the addition of new hosts as follows:

- If there is no certificate in the SAP Host Agent security directory, the SAP HANA database lifecycle manager generates one. The SAP HANA host name is used as the default certificate owner. The certificate owner can be changed by using the call option `--certificates_hostmap=<fully_qualified_domain_name>`.
- If there is an existing certificate, the following applies:
 - If the certificate host name is not passed to the SAP HANA database lifecycle manager, or if the certificate host name is the same as the owner of the current certificate, the current certificate is preserved.
 - If the certificate host name is passed via the call option `--certificates_hostmap` and it differs from the owner of the current certificate, a new certificate is generated.
 - During update of an SAP HANA system, if the certificates on all hosts are in place, the call option `--certificates_hostmap` is ignored and the current certificates are preserved.

If you want to use your own SSL certificates, see the SAP Host Agent documentation in Related Information.

Related Information

[SSL Configuration for the SAP Host Agent](#)

8.1.1.1.5.1.2.2 Starting Platform LCM Tasks as the System Administrator User <sid>adm

When starting platform LCM tasks as the system administrator user <sid>adm, the SAP HANA database lifecycle manager (HDBLCM) requires the usage of SAP Host Agent for execution of remote and local operations.

The following tasks in the SAP HANA database lifecycle manager can be performed as the system administrator user <sid>adm:

- System update from the installation medium
- Installation or update of additional components from the SAP HANA resident HDBLCM
- Host addition and host removal
- System Landscape Directory (SLD) registration configuration
- Inter-service communication configuration

Make sure that SAP Host Agent is installed and configured (HTTPS-enabled) on all hosts of the SAP HANA system.

i Note

Platform LCM tasks cannot be executed remotely via SSH as the system administrator user <sid>adm.

i Note

Make sure that the system administrator user <sid>adm has permissions to read the paths passed as parameters (for example, the locations of the SAP HANA components).

8.1.1.1.5.1.2.3 Add Hosts Using SAP Host Agent

You can add hosts to an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in combination with the SAP Host Agent in the command-line interface.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).
- The host which is to be added has access to the installation directories <sapmnt> and <sapmnt>/<SID>.
- The SAP Host Agent is installed on the host which is to be added. The SAP Host Agent will create the <sapsys> group, if it does not exist prior to installation. Make sure that the group ID of the <sapsys> group is the same on all hosts. For information about installing or updating the SAP Host Agent individually, see *Installing SAP Host Agent Manually* and *Upgrading SAP Host Agent Manually*.
- A Secure Sockets Layer (SSL) certificate is configured for the SAP Host Agent, so that the HTTPS port 1129 is accessible and the Personal Security Environment (PSE) for the server is prepared. For more

information about SSL configuration for the SAP Host Agent, see *Configuring SSL for SAP Host Agent on UNIX*.

- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on as root user or as the system administrator user `<sid>adm`.
- The difference between the system time set on the installation host and the additional host is not greater than 180 seconds.
- The operating system administrator (`<sid>adm`) user may exist on the additional host. Make sure that you have the password of the existing `<sid>adm` user, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) resident program will not modify the properties of any existing user or group.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblcml
```

By default, `<sapmnt>` is `/hana/shared`.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblcml --remote_execution=saphostagent
```

3. Select the index for the `add_hosts` action.
4. Enter the names of the hosts to be added.
5. Enter the SAP Host Agent administrator (`sapadm`) password.
6. Define additional system properties.
7. Review the summary, and select `y` to finalize the configuration.

Results

You have added one or more new hosts to an SAP HANA system. The SAP HANA system you have configured is a multiple-host system.

The new hosts have been added to the SAP HANA landscape information. The new hosts have been added to the landscape information of the system database.

This configuration task can also be performed in batch mode and using a configuration file. For more information about the available configuration methods, see *Using the SAP HANA Platform LCM Tools*.

Related Information

[Host Addition Concepts \[page 1410\]](#)

[Using the SAP HANA Platform LCM Tools \[page 921\]](#)

[Using SAP Host Agent to Execute Platform LCM Tasks \[page 939\]](#)

[Installing SAP Host Agent Manually](#)

[Updating SAP Host Agent Manually](#)

[Configuring SSL for SAP Host Agent on UNIX](#)

8.1.1.1.5.2 Decentralized Execution of Platform LCM Tasks

In some circumstances platform LCM actions must be executed on each individual host of the multiple-host system. This is referred to as **decentralized execution**.

Typically, SAP HANA platform lifecycle management actions, such as update, rename, and inter-service communication configuration, can be performed on a multiple-host system from one host. This is referred to as **centralized execution** and requires SSH or root credentials. For more information, see Centralized Execution of Platform LCM Tasks in Related Information.

In some circumstances, a secure shell (SSH) key may not be installed or root credentials are not available. In this case, the platform LCM actions must be executed on each individual host of the multiple-host system, which is also known as **decentralized execution**. For more information about decentralized execution, see SAP Note 2048681 in Related Information.

Related Information

[SAP Note 2048681](#) 

[Executing Platform LCM Tasks \[page 937\]](#)

[Centralized Execution of Platform LCM Tasks \[page 938\]](#)

8.1.1.1.6 Additional Information About Using the SAP HANA Platform LCM Tools

If you have already familiarized yourself with the way the SAP HANA database lifecycle manager (HDBLCM) works, you may be interested in additional information like log and trace files, Linux kernel parameter settings, or troubleshooting.

Related Information

[Logging \[page 944\]](#)

[Linux Kernel Parameters \[page 944\]](#)

[General Troubleshooting for the SAP HANA Platform LCM Tools \[page 946\]](#)

8.1.1.1.6.1 Logging

SAP HANA platform lifecycle management processes are logged by the system. The log files are stored in the following path:

```
/var/tmp/hdb_<SID>_<action>_<time stamp>
```

where <action> ::= install | update | addhost | uninstall | and so on.

The following log files are written while performing the action:

- <hdbcommand>.log: can be read using a text editor
- <hdbcommand>.msg: XML format for display in the installation tool with the GUI
- <hostname>_tracediff.tgz: provides a delta analysis of the original trace files, makes a detailed analysis easier

You can also view diagnostic files in the SAP HANA database explorer using the administration function. For more information, see *View Diagnostic Files in the SAP HANA Database Explorer* in the *SAP HANA Administration Guide*.

Instant Logging

If an LCM action crashes or hangs before the execution is finished, even if no LCM action trace is enabled, HDBLCM writes a trace, which has the function of a preliminary (unformatted) log file. Upon program completion, this preliminary logfile is removed and replaced by the real, formatted log file.

The environment variable `HDB_INSTALLER_TRACE_FILE=<file>` enables the trace.

The environment variable `HDBLCM_LOGDIR_COPY=<target directory>` creates a copy of the log directory.

Log Collection

If you perform platform LCM actions on multiple-host SAP HANA systems, all log files are collected to a local folder to make error analysis more convenient.

To collect log files for multiple-host SAP HANA systems, an HDBLCM action ID is passed to each sub-program (underlying LCM tool) working on a remote host. Each sub-program writes a copy of the log file in to the following directory: `<installation path>/<SID>/HDB<instance number>/<host name>/trace`

Related Information

[View Diagnostic Files in the SAP HANA Database Explorer \[page 662\]](#)

8.1.1.1.6.2 Linux Kernel Parameters

The following table describes the parameters and limits that are set by the SAP HANA database lifecycle manager (HDBLCM) during the installation or update of an SAP HANA database. The actual values may differ, depending on your system configuration.

The SAP Host Agent can automatically optimize the following Linux Kernel Parameters:

- `net.ipv4.ip_local_port_range`
- `net.ipv4.ip_local_reserved_ports`

To configure the SAP Host Agent, make sure that the `/etc/sysctl.conf` configuration does not contain any of these two parameters. Afterwards, configure the SAP Host Agent profile parameters as described in *SAP Note 401162*.

Parameter	Description	Value	Location
<code>nofile</code>	Open file descriptors per user	1048576	<code>/etc/security/limits.conf</code>
<code>fs.file-max</code>	Open file descriptors per host	20000000	<code>/etc/sysctl.conf</code>
<code>fs.aio-max-nr</code>	Maximum number of asynchronous I/O requests	18446744073709551615 (= $2^{64}-1$ = <code>ULONG_MAX</code>)	<code>/etc/sysctl.conf</code>
<code>vm.memory_failure_early_kill</code>	Method for killing processes when an uncorrected memory error occurs	1	<code>/etc/sysctl.conf</code>
<code>kernel.shmmax</code>	Maximum shared memory segment size (the default minimum value is 1 GB)	1073741824	<code>/etc/sysctl.conf</code>
<code>kernel.shmni</code>	Maximum number of shared memory segments	32768	<code>/etc/sysctl.conf</code>
<code>kernel.shmall</code>	System-wide limit of total shared memory, in 4k pages	<ul style="list-style-type: none"> • RAM \geq 35.5 TB: $(shmmax * shmni) / 65536$ • RAM $<$ 35.5 TB: $(0.9 * \langle RAM \text{ in bytes} \rangle) / 4096$ 	<code>/etc/sysctl.conf</code>
<code>net.ipv4.ip_local_port_range</code>	Lower limit of ephemeral port range	40000	<code>/etc/sysctl.conf</code>
<div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin: 10px 0;"> <p>i Note</p> <p>No changes are applied if the SAP HANA database lifecycle manager (HDBLCM) detects or installs SAP Host Agent version 7.20.162 or greater.</p> </div>			
<code>vm.max_map_count</code>	Maximum number of Virtual Memory Areas (VMAs) that a process can own	2147483647	<code>/etc/sysctl.conf</code>

Related Information

[SAP Note 2382421](#)

8.1.1.1.6.3 General Troubleshooting for the SAP HANA Platform LCM Tools

The SAP HANA database lifecycle manager (HDBLCM) is a wrapper tool that calls the underlying HDB tools to perform the platform LCM action. If something unexpected happens when using HDBLCM, and the LCM action cannot be completed, you can check the logs and separately run the affected underlying tools.

Caution

We only recommend the following underlying tools to be used for troubleshooting purposes.

Program Name	Description	Location
hdbinst	Command-line tool for installing the software	Installation media
hdbsetup	Installation tool with a graphical interface for installing or updating the software	Installation media
hdbuninst	Command-line tool for uninstalling the software and removing a host	Installation media and <installation path>/ <SID>/global/hdb/ install/bin
hdbaddhost	Command-line tool for adding a host to a system	<installation path>/ <SID>/global/hdb/ install/bin
hdbupd	Command-line tool for updating the software	Installation media
hdbrename	Command-line tool for renaming a system	<installation path>/ <SID>/global/hdb/ install/bin and /usr/sap/<SID>/SYS/ global/hdb/ install/bin

Program Name	Description	Location
hdbreg	Command-line tool for registering an SAP HANA system	<pre><installation path>/ <SID>/global/hdb/ install/bin and /usr/sap/<SID>/SYS/ global/hdb/ install/bin</pre>
hdbremovehost	Command-line tool for removing a host	<pre><installation path>/ <SID>/global/hdb/ install/bin and /usr/sap/<SID>/SYS/ global/hdb/ install/bin</pre>
hdbmodify	<p>This command line tool removes and adds remote hosts.</p> <p>Furthermore, the listen interface can be changed ('local', 'global', 'internal').</p>	<pre><installation path>/ <SID>/global/hdb/ install/bin and /usr/sap/<SID>/SYS/ global/hdb/ install/bin</pre>
hdbupprep	Command-line tool for upgrading a repository by loading delivery units into the database	<pre><installation path>/ <SID>/global/hdb/ install/bin and /usr/sap/<SID>/SYS/ global/hdb/ install/bin</pre>

8.1.1.1.6.4 Managing SAP HANA System Components

SAP HANA system components can be installed, updated, or uninstalled using the SAP HANA database lifecycle manager (HDBLCM).

The SAP HANA system is made up of the following components:

- SAP HANA mandatory components

- SAP HANA server
- SAP HANA client
- SAP HANA additional components
 - SAP HANA studio
 - Application Function Libraries (AFL and the product-specific AFLs POS, SAL, SCA, SOP, TRD, UDF)
 - SAP liveCache applications (SAP LCA or LCAPPS-Plugin)
 - SAP HANA smart data access (SDA)

i Note

To install or uninstall the Solution Manager Diagnostics Agent, use Software Provisioning Manager (SWPM). For more information about the setting up the Solution Manager Diagnostics Agent using SWPM, see SAP Note 1858920 in Related Information.

i Note

SAP LT replication configuration is a part of SL Toolset 1.0. For more information about configuring SAP LT replication, see SAP Note 1891393 in Related Information.

- SAP HANA options
 - SAP HANA dynamic tiering
 - SAP HANA streaming analytics
 - SAP HANA accelerator for SAP ASE

For more information about installing, updating, and uninstalling the SAP HANA mandatory components and SAP HANA additional components, see the *SAP HANA Server Installation and Update Guide*. For more information about installing, updating, and uninstalling the SAP HANA options, see SAP HANA option documentation in Related Information.

⚠ Caution

Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

Related Information

[SAP Note 1858920](#)

[SAP Note 1891393](#)

[Important Disclaimer for Features in SAP HANA Platform \[page 1980\]](#)

8.1.1.1.6.5 Check the Installation Using the Command-Line Interface

You can check the installation of an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in the command-line interface for troubleshooting.

Prerequisites

- You are logged in as root user.
- Any user has read and execute permissions for the directory that contains the installation medium.
- Depending on the storage solution, set the export options `rw,no_root_squash` for the installation directory.
- The operating system administrator (`<sid>adm`) user and other operating system users may exist prior to installation. Make sure that you have the password of the existing users, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) resident program will not modify the properties of any existing user or group.
- The SAP HANA system has been installed with its server software on a shared file system (export options `rw,no_root_squash`).
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblcml
```

By default, `<sapmnt>` is `/hana/shared`.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblcml --action=check_installation
```

3. Enter the required credentials.
4. Review the summary, and select `y` to finalize the configuration.

Results

The check tool outputs basic information about the configuration of the file system, system settings, permission settings, and network configuration. The checks are based on the property file stored in the following path:

```
<sapmnt>/<SID>/global/hdb/install/support/hdbcheck.xml
```

Use the generated log files as a reference in the case of troubleshooting. The log file is stored in the following path:

```
/var/tmp/hdb_<SID>_hdblcm_check_installation_<time stamp>/hdblcm.log
```

8.1.1.2 Users Created During Installation

The following users are automatically created during the installation: `<sid>adm`, `sapadm`, and `SYSTEM`.

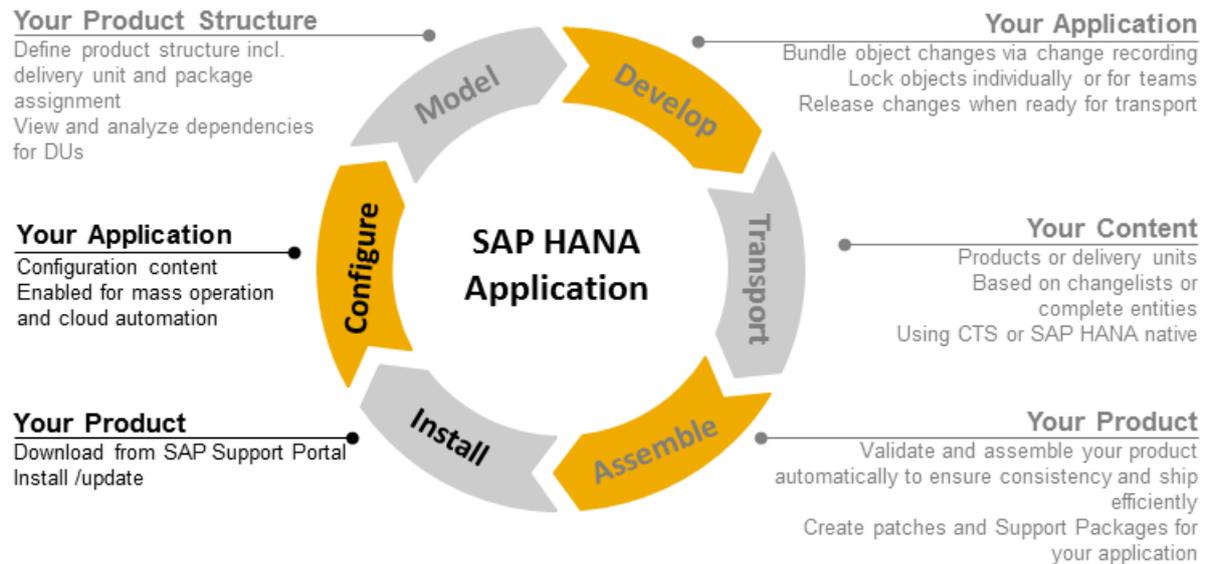
User	Description
<code><sid>adm</code>	<p>The operating system administrator.</p> <ul style="list-style-type: none">• The user <code><sid>adm</code> is the operating system user required for administrative tasks such as starting and stopping the system.• The user ID of the <code><sid>adm</code> user is defined during the system installation. The user ID and group ID of this operating system user must be unique and identical on each host of a multiple-host system.• The password of the <code><sid>adm</code> user is set during installation with the <code>password</code> parameter.
<code>sapadm</code>	<p>The SAP Host Agent administrator.</p> <ul style="list-style-type: none">• If there is no SAP Host Agent available on the installation host, it is created during the installation along with the user <code>sapadm</code>.• If the SAP Host Agent is already available on the installation host, it is not modified by the installer. The <code>sapadm</code> user and password are also not modified.• The password of the <code>sapadm</code> user is set during installation with the <code>sapadm_password</code> parameter.
<code>SYSTEM</code>	<p>The database superuser.</p> <ul style="list-style-type: none">• Initially, the <code>SYSTEM</code> user has all system permissions. Additional permissions can be granted and revoked again, however the initial permissions can never be revoked.• Two <code>SYSTEM</code> users are created: one for the system database and one for the tenant database.• The password of the <code>SYSTEM</code> user is set during installation with the <code>system_user_password</code> parameter.

Related Information

[Operating System User sidadm \[page 714\]](#)

8.2 SAP HANA Application Lifecycle Management

SAP HANA Application Lifecycle Management supports you in all phases of an SAP HANA application lifecycle, from modelling your product structure, through application development, transport, assemble, and install.



Phases of SAP HANA Application Lifecycle Management

The following are phases of SAP HANA application lifecycle management. Some phases are designed for developers only, while others, such as the installation of add-on products and software components, are designed for both.

- Model**
 You define your product structure to provide a framework for efficient software development. This includes creating the following metadata: creating Repository packages for development, defining a package hierarchy and assigning packages to delivery units. The delivery units are then bundled in products.
- Develop**
 You perform software developments in Repository packages. SAP HANA application lifecycle management supports you with change tracking functions which allow you to transport only changed objects.
- Transport**
 You can transport your developed content in different ways according to your needs. You can choose between transporting products or delivery units, based on changelists or complete entities. The transport type can be native SAP HANA transport or transport using Change and Transport System (CTS). You can also export delivery units, and import them into another system.
- Assemble**
 The developed software plus the metadata defined when modelling your product structure as well as possible translation delivery units are the basis for assembling your add-on product. You can also build Support Packages and patches for your product.

- **Install**

You can install add-on products or software components that you downloaded from SAP Support Portal or assembled yourself.

All tasks related to the **Install** and **Configure** phases of SAP HANA application lifecycle management are documented in this *SAP HANA Administration Guide*. The tasks related to software development are documented in the *SAP HANA Developer Guide (For SAP HANA Studio)*. **All** phases of SAP HANA application lifecycle management are documented in the *SAP HANA Application Lifecycle Management Guide*.

Related Information

[Installing and Updating SAP HANA Products and Software Components in SAP HANA XS Classic Model \[page 952\]](#)

8.2.1 Installing and Updating SAP HANA Products and Software Components in SAP HANA XS Classic Model

SAP HANA application lifecycle management provides functions for installing and updating SAP HANA products or individual software components of SAP HANA XS classic model that you have downloaded from the SAP Support Portal, or that you have assembled yourself.

Context

SAP HANA products consist of software components which are deployed to the SAP HANA repository. You have the following options to install and update SAP HANA products and software components:

- Using a **SAP Fiori application** integrated in the **SAP HANA Application Lifecycle Management XS application**. This application can be started in the following ways:
 - Start the SAP HANA Application Lifecycle Management on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/lm`. Afterwards, choose the **INSTALLATION** tab or tile.
 - Using a link in SAP HANA Web-based Development Workbench.
In the SAP HANA Web-based Development Workbench Editor tool, choose ► [Navigation Links](#) ► [Lifecycle Management](#) ▾. The SAP HANA Application Lifecycle Management home screen opens, where you can choose the **INSTALLATION** tab or tile..
 - Using the context menu in SAP HANA studio.
Choose ► [Lifecycle Management](#) ► [Application Lifecycle Management](#) ► [Installation](#) ▾ from the context menu for a particular system in the *SAP HANA Administration Console* perspective in SAP HANA studio.

The documentation about using SAP HANA Application Lifecycle Management to install and update SAP HANA products and software components describes the following use cases:

- *Installing and Updating SAP HANA Products*

- *Installing and Updating SAP HANA Software Components*
- Using the `hdbalm` **commandline tool**.
 To start `hdbalm`, start a command line client and navigate to the directory where `hdbalm` is located. You can also add this directory to your path.
 For more information about using `hdbalm` to install and update SAP HANA products and software components, see the following topics in the *SAP HANA Application Lifecycle Management Guide*:
 - *Using hdbalm*
 - *hdbalm install Command*

i Note

SAP HANA system components like the SAP HANA client, SAP HANA studio, and additional system components like Application Function Libraries (AFL and the product-specific AFLs POS, SAL, SCA, SOP, UDF), SAP liveCache applications (SAP LCA or LCAPPS-Plugin), XS advanced runtime applications, or SAP HANA smart data access (SDA) are installed and updated using the SAP HANA database lifecycle manager (HDBLCM). For more information, refer to the *SAP HANA Server Installation and Update Guide*.

Related Information

[Installing and Updating SAP HANA Products \[page 953\]](#)

[Installing and Updating SAP HANA Software Components \[page 955\]](#)

[Installation and Update Options \[page 958\]](#)

8.2.1.1 Installing and Updating SAP HANA Products

You can install and update SAP HANA products using SAP HANA application lifecycle management.

Prerequisites

- You have a product archive of an SAP HANA product that you want to install or update.

i Note

An SAP HANA product archive is a `*.zip` file that contains one or more software component archives as well as metadata files. For more information about the archive types that are used to deliver SAP HANA content, read the information about *SAP HANA content* in the *SAP HANA Administration Guide*.

- You have the privileges granted by a role based on the SAP HANA Application Lifecycle Management `sap.hana.xs.lm.roles::Administrator` role template.

Procedure

1. Open the SAP HANA Application Lifecycle Management.

The SAP HANA Application Lifecycle Management is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/lm`.

2. Choose the *INSTALLATION* tab.
3. Click in the *Archives* selection field to select a product archive from your file directory that you want to install or update.

The product archive is uploaded. The header area contains information about the product version, including the action that is to be performed: Installation or Update.

4. The *Product Instances* tab lists all product instances that are part of the archive.

For each product instance, the result of a software component check is displayed.

The following results can occur:

- *OK*
The product instance can be installed.
- *Downgrade*
The product instance contains one or more software components that are already installed in newer versions than the ones to be installed. The installation of this product instance would lead to a downgrade of these software components. Downgrades are not allowed. To continue the installation of the product, you have to set the installation/update option *Keep newer version of software component*. In this case, the downgrading software components will be skipped during the installation of the product instance.

Note

If it is required that you install the software component that causes the downgrade, for example, if the newer version has errors and you want to revert to the previous version, you can use the `install` command of `hdbalm` with the option `ALLOW_DU_DOWNGRADE` to enable the downgrade. However, use this option with care, since this may affect other installed products which require the newer version of this software component.

- *Some software components are installed already*
If software components are already installed in the same version, by default, the system skips their installation during the installation/update of the product instance. If you want to reinstall the same version, you can set the option *Overwrite the same version of software component* in the installation and update options.

Click in the line of the product instance to display more information about the software components that are part of the product instance. For each software component, a status is displayed, as well as the installed version and the new version. If you click on the status icon, you get more information about the status.

5. If required, set installation or update options.

The options allow you to override the default behavior of the installation or update for specific situations. Use them with care. For more information about the options, see *Installation and Update Options*.

6. Select product instances for installation.

You can individually select single product instances. To install all instances, select the *Instance* check box in the header row.

All instances of the product, that are already installed on your system will automatically be checked for updates. If the archive that you uploaded contains newer versions for one or several software components, they will automatically be updated. It doesn't matter whether you selected the respective instance for installation.

7. To start the installation, choose *Install*.

The system displays the progress of the individual installation steps. You can click on each step to expand the log of the step.

Results

If errors occur during the installation or update, an error message indicates the reason for the error and the system provides a log with more detailed information. If you cannot solve the problem and you need to open a customer message, ensure that you assign it to the message component of the SAP HANA software component or product instance that caused the error. The *Support Information* tab contains the relevant information. Do **not** assign the message to the component of SAP HANA application lifecycle management since this may slow down the problem solving process.

If the installation or update finished successfully, you can start another installation using *New Installation*.

Related Information

[Installation and Update Options \[page 958\]](#)

[Installing and Updating SAP HANA Products and Software Components in SAP HANA XS Classic Model \[page 952\]](#)

[SAP HANA Administration Guide \[page 10\]](#)

8.2.1.2 Installing and Updating SAP HANA Software Components

You can install and update SAP HANA software components using SAP HANA application lifecycle management.

Prerequisites

- You have one or multiple archives of SAP HANA software components that you want to install or update.

i Note

An SAP HANA software component archive is a *.zip file that contains one delivery unit archive (*.tgz) as well as metadata files. For more information about the archive types that are used to deliver

SAP HANA content, read the information about *SAP HANA content* in this *SAP HANA Administration Guide*.

Software components which need to be installed at the operating system level, such as Application Function Libraries (AFLs), are **not** installed using SAP HANA application lifecycle management.

- You have the privileges granted by the SAP HANA Application Lifecycle Management `sap.hana.xs.lm.roles::Administrator` role.

Procedure

1. Open the SAP HANA Application Lifecycle Management.

The SAP HANA Application Lifecycle Management is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/lm`.

2. Choose the *INSTALLATION* tab.
3. Click in the *Archives* selection field to select one or multiple software component archives from your file directory that you want to install or update.

The software component archives are uploaded. For each software component the following information is displayed:

- *Status*

The following status values exist:

- *New*

The software component is not yet installed and can be installed.

- *Update*

The software component is already installed and can be updated to a higher version.

- *Downgrade*

The software component is already installed in a newer version than the one that is to be installed. The installation of this software component would lead to a downgrade. Downgrades are not allowed. You cannot continue to install the software component.

i Note

If it is required that you install the software component that causes the downgrade, you can use the `install` command of `hdbal` with the option `ALLOW_DU_DOWNGRADE` to enable the downgrade. However, use this option with care, since this may affect already installed products which require the newer version of this software component.

- *Already installed*

If software components are already installed in the same version, by default, the system skips their installation during the installation/update. If you want to reinstall the same version, you can set the option *Overwrite the same version of software component* in the installation and update options.

- *Version* that is already *installed* in the system

- *New version* that is to be installed

- Whether the software component is part of a product that is already installed.

- *Information* relevant for the *support*

4. If required, set installation or update options.

The options allow you to override the default behavior of the installation or update for specific situations. Use them with care. For more information about the options, see *Installation and Update Options*.

5. To start the installation, choose *Install*.

The system displays the progress of the individual installation steps. You can click on each step to expand the log of the step.

Results

If errors occur during the installation or update, an error message indicates the reason for the error and the system provides a log with more detailed information. If you cannot solve the problem and you need to open a customer message, ensure that you assign it to the message component of the SAP HANA software component that caused the error. You can find this information in the support information of the component. Do **not** assign the message to the component of SAP HANA application lifecycle management since this may slow down the problem solving process.

If the installation or update finished successfully, you can start another installation using *New Installation*.

Related Information

[Installation and Update Options \[page 958\]](#)

[Installing and Updating SAP HANA Products and Software Components in SAP HANA XS Classic Model \[page 952\]](#)

[SAP HANA Content \[page 994\]](#)

8.2.1.3 Installation and Update Options

Installation and update options are available that allow you to influence the installation and update behavior, if required.

Installation and Update Options

Option	Corresponding Installation Option in hdbalm	Description
Overwrite the same version of software component	ALLOW_DU_SAME_VERSION	<p>By default, the system does not install a software component if the same version is already installed. It is possible to override this behavior in the following situations:</p> <ul style="list-style-type: none">• If a previous installation operation failed, for example, because of activation errors.• If you run continuous integration scenarios in which the same version of a software component is installed regularly.
Keep newer version of software component	ALLOW_KEEP_DU_NEWER_VERSION	<p>If the product instance to be installed contains software components with lower versions than the installed ones, installing the software component in the lower version would lead to a downgrade of this software component. This is not allowed. You cannot install a downgrading software component. If you want to install the product instance without the downgrading software components, you can use this option. This option is useful if a software component is part of several products. If the product to be installed contains the software component in a version which is lower than the one already installed, you can choose to retain the newer version. In this case, the installation of the software component is skipped.</p>
Allow version update	ALLOW_DU_VERSION_UPDATE	<p>Allows version updates of software components.</p> <p>In some cases, for example, if a software component is part of several products, a version update of a software component could render one product inoperable. If the system detects an inconsistency, it aborts the operation. You can use this option to turn off this behavior.</p>

Option	Corresponding Installation Option in hdbalm	Description
Roll back installation if activation errors occur (default)	This is the default behavior in hdbalm.	<p>By default, the installation is canceled if any activation errors occur and the complete installation is rolled back.</p> <p>Installation is also rolled back if you modified objects in your system and a modified object cannot be activated because it references an object that is part of the installation archive. This can occur, for example, if a procedure or view references a table in the archive.</p>
Ignore activation errors of referencing objects	USE_TWO_COMMIT_ACTIVATION	<p>If an installation fails since an object outside of the archive cannot be activated due to references to an object in the archive, you can repeat the installation with this activation option. In this case, the object remains broken in the system after the installation, but the installation itself finishes successfully. You must correct the errors manually after the installation.</p> <p>You can check the transport log after performing the installation without this option to find out whether the activation errors were caused by objects in the archive or outside of the archive. After repeating the installation with this option, check the transport log to find out which objects must be repaired afterwards.</p>

Related Information

[Installing and Updating SAP HANA Products and Software Components in SAP HANA XS Classic Model \[page 952\]](#)

8.2.2 Installing and Updating Products and Software Components in SAP HANA XS Advanced Model

Application lifecycle management for SAP HANA XS advanced model provides functions for installing and updating products as well as individual software components of SAP HANA XS advanced that you have downloaded from the SAP Support Portal.

Prerequisites

- The prerequisites described under *Prerequisites and Authorizations* are fulfilled. The link to the topic is in the *Related Information*.
- You have an SAP HANA XS advanced product or software component archive that you want to install or update.

i Note

An SAP HANA XS advanced software component archive is a *.zip file that consists of a multi-target application archive (MTA archive = *.mtar file) and an `SL_MANIFEST.xml` file that contains metadata, such as version, vendor, support package, and patch level for the MTA archive.

An SAP HANA XS advanced product archive is a *.zip file that consists of one or multiple software component archives plus a `pd.xml` and a `stack.xml` file. Both files contain metadata for the product, such as version, support package level, and vendor.

Context

i Note

From SPS 11, SAP HANA includes an additional run-time environment for application development: SAP HANA extended application services (XS), advanced model. SAP HANA XS advanced model represents an evolution of the application server architecture within SAP HANA by building upon the strengths (and expanding the scope) of SAP HANA extended application services (XS), classic model. SAP recommends that customers and partners who want to develop new applications use SAP HANA XS advanced model. If you want to migrate existing XS classic applications to run in the new XS advanced run-time environment, SAP recommends that you first check the features available with the installed version of XS advanced; if the XS advanced features match the requirements of the XS classic application you want to migrate, then you can start the migration process.

You have the following options to install and update SAP HANA products and software components in SAP HANA XS advanced:

- Using the XS advanced command line interface (CLI)
- Using the XS Advanced Application Lifecycle Management graphical user interface

The links to the corresponding topics are in the *Related Information*.

i Note

SAP HANA system components like the SAP HANA client, SAP HANA studio, and additional system components like Application Function Libraries (AFL and the product-specific AFLs POS, SAL, SCA, SOP, UDF), SAP liveCache applications (SAP LCA or LCAPPS-Plugin), XS advanced runtime applications, or SAP HANA smart data access (SDA) are installed and updated using the SAP HANA database lifecycle manager (HDBLCM). For more information, refer to the *SAP HANA Server Installation and Update Guide*.

Related Information

[Prerequisites and Authorizations \[page 961\]](#)

[Installing and Updating Using the Command Line Interface \[page 962\]](#)

[Installing and Updating Using the XS Advanced Application Lifecycle Management Graphical User Interface \[page 981\]](#)

8.2.2.1 Prerequisites and Authorizations

The following prerequisites have to be fulfilled when you use functions required for installing and updating SAP HANA products and software components in SAP HANA XS advanced model.

- The XS advanced run time is installed and available on the SAP HANA server.
For more information, see *Installing XS Advanced Runtime* in the *SAP HANA Server Installation and Update Guide*.
- When using the XS Advanced Application Lifecycle Management Graphical User Interface, the following software components are installed in addition to the XS advanced run time:
 - SAP UI5 component: XSACUI5FESV344P in version 1.44.8 or higher
 - SAP HANA XS Advanced Application Lifecycle Management Product Installer UI:
XSAC_ALM_PRODUCT_INSTALLER_UI1

You can download the components from the Software Download Center at <https://support.sap.com/swdc> and install them using the `xs install` command in the command line interface.

For more information, see the *SAP HANA Server Installation and Update Guide*.

i Note

If you have performed the default SAP HANA medium installation, both the SAP UI5 component and the SAP HANA XS Advanced Application Lifecycle Management Product Installer UI are already installed.

- Optional when using the XS advanced command line interface (XS CLI) for installation and update: The XS advanced command line client is installed on your local machine.
The XS CLI client tools are installed by default on the SAP HANA server. You can log on to the server and execute the installation command there. However, if you want to connect to SAP HANA from your local machine, you must download and install the client tools locally. The XS CLI client tools (`xs.onpremise.runtime.client_<platform>-<version>.zip`) can be downloaded from the SAP HANA server, from the installation DVD, or from the SAP support portal.

- The SAP HANA database user that is used to perform the installation or update has one of the following permissions assigned:
 - The user has the `XS_CONTROLLER_USER` parameter assigned as well as the *SpaceDeveloper* role for each space in which the user wants to perform an installation or update.
 - The user has the `XS_CONTROLLER_ADMIN` parameter assigned.
This scope allows the installation in all spaces.

For more information on assigning roles in SAP HANA XS advanced, see *Setting Up Security Artifacts* in the *SAP HANA Administration Guide*.

Related Information

[SAP HANA Administration Guide \[page 10\]](#)

8.2.2.2 Installing and Updating Using the Command Line Interface

To install and update products and software components in SAP HANA XS advanced, the `xs install` command is available in the XS advanced command line interface (CLI). Using this command you can install or update one product archive or one software component archive at a time.

Procedure

1. Start the XS advanced command line interface (CLI).
2. Log on to the SAP HANA XS advanced runtime in the organization and space in which you want to install or update the product or software component.

To do this, use the `xs login` command with the following arguments and options:

Argument/Option	Description
<code>-u</code>	SAP HANA database user with the permissions as described in the <i>Prerequisites</i> section
<code>-p</code>	password
<code>-o</code>	organization in which the installation or update takes place
<code>-s</code>	space in which the installation or update takes place

Sample Code

```
xs login -u demo -p test -o demoorg -s demospace
```

For more information, see the *XS CLI: Logon and Setup* topic in *SAP HANA Developer Guide (for SAP HANA XS Advanced Model)*.

3. If you want that the archives to be installed are checked with antivirus software before the installation or update process, proceed as described under *Set Up a Virus Scan for Installation Archives*. The link to the topic is in the *Related Information* section.
4. Start the installation or update of the product or software component.

The `xs install` command is available in the XS advanced CLI both for installing product and software component archives in XS advanced and updating these. The `xs install` command detects whether the archive is a product archive or a software component archive. It also detects whether the product or software component is installed already and subsequently executes either an installation or update operation.

Enter the `xs install` command and specify the path to the archive. If required, enter any additional options. For example, to install a specific instance of a product, you can use the `-i` option and specify the product instance. Or to make sure that the entity you are about to install is a product, you can add the `-pv` option. In this case, the installation is only performed if you specify a product archive for the `xs install` command. If you specify a software component archive, the installation is not performed.

Sample Code

```
xs install ../sap_demo/target/XSASAMPLEPRODUCT1.0.zip
```

Instead of `xs install` you can also use the `xs ins` alias. For more information on the options, see *Installation and Update Options in XS Advanced Model*. For installation examples, see *Examples: Installing and Updating Products and Software Components in XS Advanced Model*. The links are in the *Related Information* section.

Results

Before installing or updating the product or software component, the system performs different checks. If no errors are found, the system performs the installation or update with the arguments and options that you specified. During the process, the product installer calls the deploy service that performs the actual deployment. Afterwards, the product installer registers the product or software component as installed.

If the installation or update cannot be performed, it is possible, in some situations, to use additional options to override the default behavior of the system. For more information, see *Checks Before Installing or Updating Products or Software Components in XS Advanced Model* and *Installation and Update Options in XS Advanced Model*.

If errors occur during the installation or update, an error message indicates the reason for the error and the system provides a log with more detailed information. If you cannot solve the problem and you need to open a customer message, ensure that you assign it to the message component of the SAP HANA product or software

component that caused the error. Do **not** assign the message to the component of SAP HANA application lifecycle management since this may slow down the problem solving process.

To display the correct log file, use one of the following commands with the process ID that you find in the result of the installation or update process.

- To display the log of a product installation, use the `display-installation-logs` command with the `--pv` option.

```
xs display-installation-logs <process ID> --pv
```

- To display the log of a software component installation, use the `display-installation-logs` command with the `--scv` option.

```
xs display-installation-logs <process ID> --scv
```

To display the history of installation or uninstallation processes, you can use the `display-installation-history` command.

```
xs display-installation-history
```

For more information on the commands used for installation, use `xs help <command>` in the XS advanced CLI, or see *The XS Command-Line Interface Reference* section in the *SAP HANA Developer Guide for XS Advanced Model*.

Related Information

[Installation and Update Options in XS Advanced Model \[page 968\]](#)

[Examples: Installing and Updating Products and Software Components in XS Advanced Model \[page 971\]](#)

[Checks Before Installing or Updating Products or Software Components in SAP HANA XS Advanced Model \[page 966\]](#)

[Set Up a Virus Scan for Installation Archives \[page 964\]](#)

8.2.2.2.1 Set Up a Virus Scan for Installation Archives

You can set an environment variable in your system to enable a default virus scan for all software component archives that you want to install or update.

Prerequisites

You have installed and configured the SAP virus scan interface as described in SAP Note [786179](#).

Context

If the antivirus software that you are using does not check the software component archives that you want to install or update, you can use the SAP virus scan interface and set the environment variable `SCAN_UPLOADS` to the value `true`. This way, the system checks all archives that you want to install or update.

By default, no antivirus protection is set for the product installer.

Procedure

1. In the commandline tool, set the XS advanced environment variable `SCAN_UPLOADS` to `true`.

Sample Code

```
xs set-env product-installer SCAN_UPLOADS true
```

For more information about setting environment variables in XS advanced, see *XS CLI: Application Management* in the *SAP HANA Developer Guide For SAP HANA XS Advanced Model*.

2. Restart the product installer.

The restart is required to ensure that the change to the environment variable takes effect.

Sample Code

```
xs restart product-installer
```

For more information about restarting applications in XS advanced, see *XS CLI: Application Management* in the *SAP HANA Developer Guide For SAP HANA XS Advanced Model*.

Related Information

[Installing and Updating Products and Software Components in SAP HANA XS Advanced Model \[page 960\]](#)

8.2.2.2.2 Checks Before Installing or Updating Products or Software Components in SAP HANA XS Advanced Model

To ensure consistency of SAP HANA products, the system executes different checks before installing or updating a product or a software component in SAP HANA XS advanced.

Product installations only: Check whether the product to be installed is already installed and in which version

If the product to be installed is not yet installed, the installation will be performed. If it is already installed, the system checks the installed version. If it is already installed in the same version, or in a lower support package level, the installation or update will be performed.

- Product is already installed in higher version
If the version of the product to be installed is lower than the installed version, the system terminates the process because installing the lower version would lead to a downgrade of the product.
You can override this behavior and allow a downgrade of the product. To do this, you can use the `ALLOW_PV_DOWNGRADE` option with the `xs install` command.
- Product is already installed in lower version
If the version of the product to be installed is higher than the installed version, the system updates the installed version automatically.

i Note

The version of a product usually consists of one or more numbers in an ascending order. In addition to the version number, a support package level is provided for the product.

Example

The version number is 1.0. In this case, the following versions are considered version updates: 1.1, 2.0, or 2.

Check whether the software component is already installed and in which version

If the software component to be installed was not yet installed, the installation will be performed. If it was already installed, the system checks the installed version. If it is installed in a lower support package or patch level, the update will be performed.

i Note

The version of a software component has the following form: "#.#.#", for example 1.0.3, where

- 1 = the version
- 0 = the support package level
- 3 = the patch level

- Software component is already installed in higher version
If a version of a software component to be installed is lower than an installed version, the system terminates the installation.
You have the following options to override this behavior:
 - You can allow a downgrade of the software component. To do this, use the `ALLOW_SC_DOWNGRADE` option.
 - For product installation only: You can skip the installation of all software components that are part of the archive and that are already installed in higher versions. To do this, use the `ALLOW_KEEP_SC_NEWER_VERSION` option.
- Software component is already installed in same version
If a version of a software component to be installed is the same as the installed version, the system proceeds as follows:
 - Product installation: The system does not install this software component. The installation of this software component is skipped during the installation of the product.
 - Software component installation: The system terminates the installation.You can override this behavior and allow the reinstallation of the same version. To do this, use the `ALLOW_SC_SAME_VERSION` option for this software component.

i Note

If the software component is installed in the system in the same version with the status *BROKEN*, it is automatically reinstalled.

- Software component is already installed in lower version
If a version of a software component to be installed is higher than an installed version, the system updates the installed version automatically.

Check for dependencies on SAP HANA platform components or other XS advanced components

If the software component has dependencies on SAP HANA platform components or other XS advanced components that are not installed, the system terminates the process and displays the missing software components. You must install or update the missing software components before you can restart the current installation or update.

For more information on the options to override the default behavior, see *Installation and Update Options in XS Advanced Model*. The link is in the *Related Information* section.

Check whether extension descriptor is valid, if an extension descriptor is used

If an extension descriptor is used for the installation process, the system checks that the extension descriptor file does not exceed a specific file size and that the syntax of the extension descriptor file is correct. If the file is too big or if the syntax is incorrect, the system will not start the installation process.

For more information on extension descriptors, see the *The MTA Deployment Extension Description* topic in the *SAP HANA Developer Guide for SAP HANA XS Advanced Model*.

Related Information

[Installing and Updating Products and Software Components in SAP HANA XS Advanced Model \[page 960\]](#)

[Installation and Update Options in XS Advanced Model \[page 968\]](#)

[Display installed Products and Software Components in XS Advanced Model \[page 976\]](#)

8.2.2.2.3 Installation and Update Options in XS Advanced Model

Installation and update options are available in SAP HANA XS advanced that allow you to influence the installation and update behavior, if required.

The following is the default syntax for the `xs install` command in the XS advanced CLI:

```
xs install <ARCHIVE> [-p <TARGET_PLATFORM>] [-pv | --PRODUCT_VERSION] [-scv | --SOFTWARE_COMPONENT_VERSION] [-t <TIMEOUT>] [-e <EXT_DESCRIPTOR_1>[,<EXT_DESCRIPTOR_2>]] [-o <VERSION_OPTION_1>[,<VERSION_OPTION_2>]] [-i | --INSTANCES <INSTANCE_1>[,<INSTANCE_2>]] [--delete-services] [--delete-service-brokers] [--no-start] [--ignore-lock]
```

The following is an example of a product installation:

```
xs install /sap_demo/target/XSASAMPLEPRODUCT1.0.zip -pv -o ALLOW_SC_SAME_VERSION
```

For more installation examples, see *Examples: Installing and Updating Products and Software Components in XS Advanced Model*. The link is in the *Related Information* section.

Installation and Update Arguments

Argument	Description
<ARCHIVE>	The path to (and name of) the archive containing the product or software component (SCV) to install, update, or downgrade

Installation and Update Options

Option	Description
-p <TARGET_PLATFORM>	Specify the target platform where the product or software component will be installed. If not specified explicitly, a target platform is created implicitly as '<ORG> <SPACE>'.
-pv --PRODUCT_VERSION	Install a product. The installation is performed only if the given archive is a product archive. Otherwise, the installation will fail.
-scv -- SOFTWARE_COMPONENT_VERSION	Install a software component. The installation is performed only if the given archive is a software component archive. Otherwise, the installation will fail.
-e <EXT_DESCRIPTOR_1>[, <EXT_DESCRIPTOR_2>]	Define one or more extensions to the installation/deployment descriptors; multiple extension descriptors must be separated by commas. For more information on extension descriptors, see the <i>The MTA Deployment Extension Description</i> topic in the <i>SAP HANA Developer Guide for SAP HANA XS Advanced Model</i> and <i>The Multi-Target Application Model</i> guide.
-t <TIMEOUT>	Specify the maximum amount of time (in seconds) that the installation service must wait for the installation operation to complete

Option	Description
<code>-o</code> <code><VERSION_OPTION_1>[,<VERSION_OPTION_2>]</code>	<p>Specify options which can be used to override the default behavior of the <code>install</code> command. The following options are available:</p> <ul style="list-style-type: none"> ALLOW_PV_DOWNGRADE Allows a downgrade of the product. By default, the system does not install a product if the product is already installed in a higher product version or support package stack since this would lead to a downgrade of the product. It is possible to override this behavior, for example, if the newer version has errors and you want to revert to the previous version. This option is available for product installations only. ALLOW_KEEP_SC_NEWER_VERSION Skips the installation of a software component if a newer version is already installed in the system. By default, the system does not install a product if a newer version of one of the software components contained in the product archive is already installed. It is possible to override this behavior. This option is useful, for example, if a software component is part of several products. If the product to be installed contains the software component in a lower version than the one already installed, you can choose to retain the newer version. If you use this option, the installation of this software component is skipped. This option is available for product installations only. ALLOW_SC_DOWNGRADE Allows a downgrade of the software component. By default, the system does not install a software component if this leads to a downgrade of the software component. It is possible to override this behavior, for example, if the newer version has errors and you want to revert to the previous version. <div data-bbox="719 1375 1394 1480" style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 5px;"> <p>⚠ Caution Use this option carefully.</p> </div> <ul style="list-style-type: none"> ALLOW_SC_SAME_VERSION Reinstalls the same version of the software component. By default, the system does not install a software component, if the same version is already installed. It is possible to override this behavior, for example, if you run continuous integration scenarios in which the same version of a software component is installed regularly.
<code>-i, --INSTANCES</code> <code><INSTANCE_1>[,<INSTANCE_2>]</code>	<p>By default all instances are installed; a comma-separated list of instances can be specified to limit the number of instances installed</p>
<code>--delete-services</code>	<p>Recreate changed services and/or delete discontinued services</p>
<code>--delete-service-brokers</code>	<p>Delete discontinued service brokers</p>

Option	Description
<code>--no-start</code>	Do not start applications that are updated during the installation
<code>--ignore-lock</code>	Force installation even if the space targeted for installation is locked
<code>--deploy-passthrough</code> <{"<key>" : "<value>", ...}>	Option for the deploy service

See also the `xs install` command reference in the *XS CLI: Plugins* topic in the *SAP HANA Developer Guide for SAP HANA XS Advanced Model*.

Related Information

[Installing and Updating Products and Software Components in SAP HANA XS Advanced Model \[page 960\]](#)

[Examples: Installing and Updating Products and Software Components in XS Advanced Model \[page 971\]](#)

[The Multi-Target Application Model](#)

8.2.2.2.4 Examples: Installing and Updating Products and Software Components in XS Advanced Model

The examples show how you can use the `xs install` command.

In the following examples you must be logged on to the XS command line interface (XS CLI) with a user with the authorizations required for installation and in the organization and space in which you want to perform the installation.

Installation of New Product

The following example installs the product *XSA Sample Product* in version 1.0, SPS 0 (initial shipment stack) contained in the file `XSASAMPLEPRODUCT_1.0.zip`:

```
XSA Sample Product (sap.com) 1.0 SPS 0
Product Instance 1
SCV_A 1.0.0
SCV_B 1.0.0
```

No version of this product is currently installed. The following command is used:

```
xs install XSASAMPLEPRODUCT_1.0.zip
```

After the installation, the `xs list-products` command displays the product as follows:

name	vendor	version	SPS	instance ids
XSA Sample Product	sap.com	1.0	0	1

The detail display for *XSA Sample Product* looks as follows:

```
xs list-products "XSA Sample Product"
-----
name                XSA Sample Product
vendor              sap.com
version             1.0
support package stack 0
-----
instance id        software component    version    state
-----
1                  -                    -          INSTALLED
                  SCV_A                1.0.0     INSTALLED
                  SCV_B                1.0.0     INSTALLED
-----
```

Update with Support Package Stack

The following example installs the product *XSA Sample Product* in version 1.0, SPS 5 contained in the file XSASAMPLEPRODUCT_1.0.5.zip in the system:

```
XSA Sample Product (sap.com) 1.0 SPS 5
Product Instance 1
SCV_A 1.5.0
SCV_B 1.5.0
```

Version 1.0, SPS 0 (initial shipment stack) of *XSA Sample Product* containing software components SCV_A in version 1.0.0 and SCV_B in version 1.0.0 is currently installed. To ensure that the archive to be installed is a product archive, the `-pv` option is used.

```
xs install XSASAMPLEPRODUCT_1.0.5.zip -pv
```

After the update, the `xs list-products` command displays the product as follows:

```
xs list-products "XSA Sample Product"
-----
name                XSA Sample Product
vendor              sap.com
version             1.0
support package stack 5
-----
instance id        software component    version    state
-----
1                  -                    -          INSTALLED
                  SCV_A                1.5.0     INSTALLED
                  SCV_B                1.5.0     INSTALLED
-----
```

Installation of Lower Support Package Version

The following example installs the product *XSA Sample Product* in version 1.0, SPS 3 contained in the file `XSASAMPLEPRODUCT_1.0.3.zip` in the system:

```
XSA Sample Product (sap.com) 1.0 SPS 3
Product Instance 1
SCV_A 1.3.0
SCV_B 1.3.0
```

Version 1.0, SPS 5 of *XSA Sample Product* containing software components `SCV_A` in version 1.5.0 and `SCV_B` in version 1.5.0 is currently installed. If the installation was started without any options, it would fail. To allow the downgrade of the support package version, you must use the `ALLOW_PV_DOWNGRADE` option. In addition, to allow a downgrade of the software components, you must use the `ALLOW_SC_DOWNGRADE` option.

```
xs install XSASAMPLEPRODUCT_1.0.3.zip -o ALLOW_PV_DOWNGRADE, ALLOW_SC_DOWNGRADE
```

After the installation, the `xs list-products` command displays the product as follows:

```
xs list-products "XSA Sample Product"
-----
name                XSA Sample Product
vendor              sap.com
version             1.0
support package stack 3
-----
instance id        software component    version    state
-----
1                  -                    -          INSTALLED
                  SCV_A                1.3.0     INSTALLED
                  SCV_B                1.3.0     INSTALLED
-----
```

Installation of Higher Product Version

The following example installs the product *XSA Sample Product* in version 2.0, SPS 1 contained in the file `XSASAMPLEPRODUCT_2.0.1.zip` in the system:

```
XSA Sample Product (sap.com) 2.0 SPS 1
Product Instance 1
SCV_A 2.1.0
SCV_B 2.1.0
```

Version 1.0, SPS 3 of *XSA Sample Product* containing software components `SCV_A` in version 1.3.0 and `SCV_B` in version 1.3.0 is currently installed. The installation of version 2.0, SPS 1 of the *XSA Sample Product* updates both the product version and the software component versions automatically.

```
xs install XSASAMPLEPRODUCT_2.0.1.zip -pv
```

After the installation, the `xs list-products` command displays the product as follows:

```
xs list-products "XSA Sample Product"
-----
name                XSA Sample Product
```

```

vendor                sap.com
version               2.0
support package stack 1
-----
instance id          software component    version    state
-----
1                    -                    -          INSTALLED
                   SCV_A                2.1.0     INSTALLED
                   SCV_B                2.1.0     INSTALLED
-----

```

Installation of Lower Product Version

The following example installs the product *XSA Sample Product* in version 1.5, SPS 3 contained in the file `XSASAMPLEPRODUCT_1.5.3.zip` in the system:

```

XSA Sample Product (sap.com) 1.5 SPS 3
Product Instance 1
SCV_A 1.3.5
SCV_B 1.3.5

```

Version 2.0, SPS 1 of *XSA Sample Product* containing software components `SCV_A` in version 2.1.0 and `SCV_B` in version 2.1.0 is currently installed. To allow a downgrade of the product version, you must use the `ALLOW_PV_DOWNGRADE` option with the command. In addition, to allow a downgrade of the software components, you must use the `ALLOW_SC_DOWNGRADE` option.

```

xs install XSASAMPLEPRODUCT_1.5.3.zip -o ALLOW_PV_DOWNGRADE, ALLOW_SC_DOWNGRADE

```

After the installation, the `xs list-products` command displays the product as follows:

```

xs list-products "XSA Sample Product"
-----
name                XSA Sample Product
vendor              sap.com
version             1.5
support package stack 3
-----
instance id          software component    version    state
-----
1                    -                    -          INSTALLED
                   SCV_A                1.3.5     INSTALLED
                   SCV_B                1.3.5     INSTALLED
-----

```

Installation of Software Component

The following example installs the software component `SCV_A` in version 1.2.3 contained in the file `SCV_A_123.zip`. No version of this software component is currently installed. The `-scv` option is used to make sure that the archive to be installed is a software component archive.

```

xs install SCV_A_123.zip -scv

```

After the installation, the `xs list-components` command displays the software component as follows:

```
xs list-components
software component          version
-----
SCV_A (sap.com)            1.2.3
```

Installation of Product with Lower Version of Software Component

The following example installs the product *XSA Test Product* in version 1.0, `SPS 3` contained in the file `XSATESTPRODUCT_1.0.3.zip`. No version of the product is currently installed. However, the product contains the software component `SCV_A` in version 1.0.3 which was already installed individually in version 1.2.3.

You have the following options to proceed with the installation:

- To allow a downgrade of the software component, you can use the `ALLOW_SC_DOWNGRADE` option with the command.

```
xs install XSATESTPRODUCT_1.0.3.zip -o ALLOW_SC_DOWNGRADE
```

After the installation, the `xs list-components` command displays the software component as follows:

```
xs list-components
software component          version
-----
SCV_A (sap.com)            1.0.3
```

- To keep the newer version of the software component, you can use the `ALLOW_KEEP_SC_NEWER_VERSION` option with the command.

```
xs install XSATESTPRODUCT_1.0.3.zip -o ALLOW_KEEP_SC_NEWER_VERSION
```

After the installation, the `xs list-components` command displays the software component as follows:

```
xs list-components
software component          version
-----
SCV_A (sap.com)            1.2.3
```

8.2.2.2.5 Display installed Products and Software Components in XS Advanced Model

To display products and software components of SAP HANA XS advanced that are already installed, the `xs list-products` and `xs list-components` commands are available.

Prerequisites

The prerequisites are fulfilled as described in the *Prerequisites and Authorizations* topic. The link is in the *Related Information* section.

Context

Instead of `xs list-products` you can also use the `xs lp` alias. Instead of `xs list-components` you can also use the `xs lc` alias.

For more information, see the *XS CLI: Plugins* topic in *SAP HANA Developer Guide (for SAP HANA XS Advanced Model)*.

Procedure

1. Start the XS advanced command-line interface (CLI).
2. Log on to the SAP HANA XS advanced runtime in the organization and space where you want to display installed products or software components.
3. You have the following options:
 - To display all products that are installed in the current organization and space, use the `xs list-products` command without any arguments.

```
xs list-products
```

The system lists all installed products with information about vendor, version, support package level and installed instances.

- To display all software components that are installed in the current organization and space, use the `xs list-components` command.

```
xs list-components [--all]
```

The system lists all installed software components with information about vendor and version. The version is displayed in the format `<software component version>.<support package level>.<patch level>`.

If you use the `--all` option with the `list-components` command, the system also displays software components for which installations have failed and which are in status *BROKEN*.

- To display detailed information for a specific installed product, use the `xs list-products` command and add the name of the product `<PRODUCT NAME>` as argument. Optionally, or if another product with the same name and different vendor exists, add the `<VENDOR>`.

❁ Example

```
xs list-products XSASAMPLEPRODUCT sap.com
```

i Note

If the product name contains a space, enter the product name in quotation marks: `xs list-products "SAMPLE PRODUCT" sap.com`

The system lists the specified product with information about vendor, version, and support package level. In addition, it lists all installed product instances and the software components that are assigned to the instances. For these, it lists the version and the state in which the software component exists in the system. They can have the following states:

- **INSTALLED**: The software component or product instance is successfully installed.
- **BROKEN**: The software component is installed in a broken state. This status can occur, for example, if there was an error in the deploy step during installation.
- **INSTALLING**: The installation of this software component is currently running.
- **INCOMPLETE**: The installation of this product instance is incomplete.
- **MISSING**: This software component is missing.

❁ Example

The output for the `xs list-products` command can look as follows:

≡ Sample Code

```
name          vendor      version  SPS   instance ids
-----
XSA Sample Product  sap.com   1.0     0     1,3
```

The output for the `xs list-products "XSA Sample Product"` command can look as follows:

≡ Sample Code

```
-----
name          XSA Sample Product
vendor        sap.com
version       1.0
SP            0
-----
instance id   software component      version  state
-----
1             -                        1.0     INSTALLED
             JAVA_HELLO_XSA_B        1.0.0   INSTALLED
             JAVA_HELLO_XSA_A        1.0.0   INSTALLED
3             -                        1.0     INSTALLED
             JAVA_HELLO_XSA_C        1.1.0   INSTALLED
             JAVA_HELLO_XSA_D        1.1.0   INSTALLED
-----
```

Related Information

[Prerequisites and Authorizations \[page 961\]](#)

8.2.2.2.6 Uninstall Products and Software Components in SAP HANA XS Advanced Model

Application lifecycle management for SAP HANA XS advanced model provides functions for uninstalling products as well as individual software components of SAP HANA XS advanced.

Prerequisites

- The prerequisites described under *Prerequisites and Authorizations* are fulfilled. The link to the topic is in the *Related Information*.
- You have a product or software component of SAP HANA XS advanced that you want to remove.

Context

You can uninstall products and software components of SAP HANA XS advanced that were installed using the `xs install` command.

Procedure

1. Start the XS advanced command-line interface (CLI).
2. Log on to the SAP HANA XS advanced runtime in the organization and space where you want to uninstall an installed product or software component.
3. Optional: Display the product or software component using the `xs list-products` or `xs list-components` command.
4. Start the uninstallation of the product or software component.

Enter the `xs uninstall` command and specify the name of the product or software component to be uninstalled, as well as the vendor, if required. In addition, you can enter options as required. The following is the default syntax for the `xs uninstall` command in the XS advanced CLI:

```
xs uninstall <NAME> [<VENDOR>] [-pv | --PRODUCT_VERSION] [-scv | --  
SOFTWARE_COMPONENT_VERSION] [-f] [--ignore-scv-reuse] [--delete-services] [--  
delete-service-brokers] [--ignore-lock]
```

The following arguments are available:

Uninstallation Arguments

Uninstallation Argument	Description
<NAME>	The name of an installed product version (PV) or software component version (SCV)
[<VENDOR>]	The name of the vendor of the specified product or software component version; optional : only needed when the same product or software component name exists with different vendors

The following options are available:

Uninstallation Options

Uninstallation Option	Description
-pv --PRODUCT_VERSION	Remove the specified product. To make sure that the entity you are about to uninstall is a product, you can add the -pv option. In this case, the uninstallation is only performed if you specify a product name as <NAME>. If you specify a software component name, the uninstallation will fail.
-scv --SOFTWARE_COMPONENT_VERSION	Remove the specified software component. To make sure that the entity you are about to uninstall is a software component, you can add the -scv option. In this case, the uninstallation is only performed if you specify a software component name as <NAME>. If you specify a product name, the uninstallation will fail.
--ignore-scv-reuse	Remove the specified software component even if it is used in other installed products. i Note You can use this option for uninstalling software components only. By default, a software component will not be uninstalled if it is also part of another installed product. You can override this behavior by using the --ignore-scv-reuse option.
-f	Remove the specified product or software component without any system prompts or confirmation

Uninstallation Option	Description
<code>-i, --INSTANCES</code> <code><INSTANCE_1>[, <INSTANCE_2>]</code>	By default all instances are uninstalled; a comma-separated list of instances can be specified to limit the number of instances to be uninstalled
<code>--delete-services</code>	Recreate changed services and/or delete discontinued services
<code>--delete-service-brokers</code>	Delete discontinued service brokers
<code>--ignore-lock</code>	Force removal of the product or software component even if the target space is locked

Sample Code

```
xs uninstall 'XSA SAMPLE PRODUCT' -pv
```

Instead of `xs uninstall` you can also use the `xs uninst` alias. For more information on the `xs uninstall` command, use the `xs help uninstall` command.

Results

The system undeploys and unregisters the specified product or software component from the SAP HANA server in the organization and space to which you are logged on.

If errors occur during the uninstallation, an error message indicates the reason for the error and the system provides a log with more detailed information. If you cannot solve the problem and you need to open a customer message, ensure that you assign it to the message component of the SAP HANA product or software component that caused the error. Do **not** assign the message to the component of SAP HANA application lifecycle management since this may slow down the problem solving process.

To display the correct log file, use the `xs display-installation-logs` command with the log ID that you find in the result of the uninstallation process and one of the `--unins_scv` or `--unins_pv` options.

```
xs display-installation-logs <log ID> --unins_scv
```

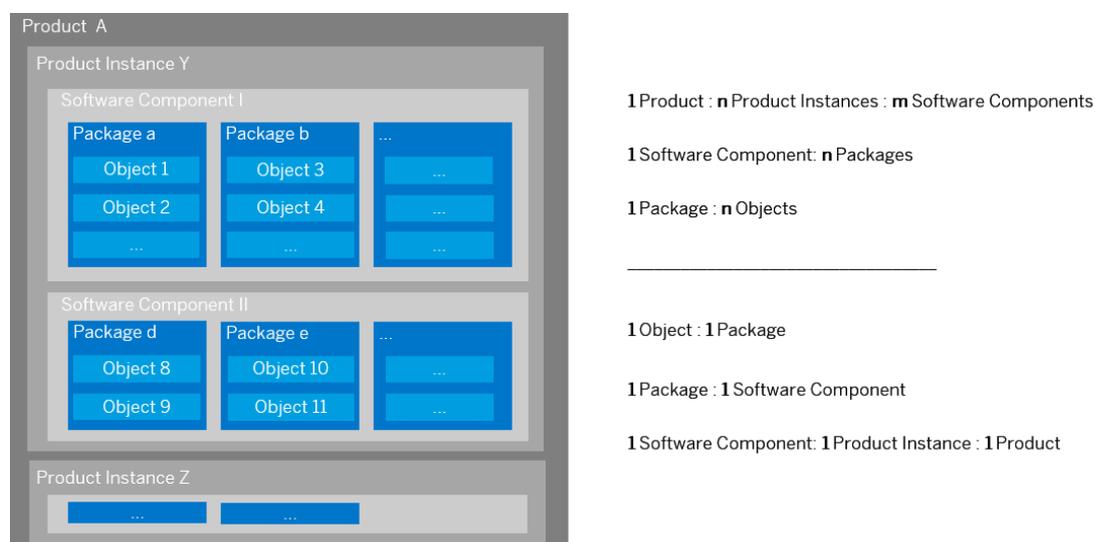
Related Information

[Prerequisites and Authorizations \[page 961\]](#)

8.2.2.3 Installing and Updating Using the XS Advanced Application Lifecycle Management Graphical User Interface

You can use the XS Advanced Application Lifecycle Management graphical user interface to install, update, and uninstall products and software components in SAP HANA XS advanced model.

An SAP HANA product consists of one or several software components and can have one or several instances. Some SAP HANA software components require an MTA extension descriptor when they are installed. If this is the case for your product or software component, you can upload one or more MTA extension descriptor files together with the installation file. For more information on MTA extension descriptors, see *The MTA Deployment Extension Description* in the *SAP HANA Developer Guide for XS Advanced Model* and *The Multi-Target Application Model* guide. The links can be found in the *Related Information* section.



Structure of an SAP HANA Product in XS Advanced

Accessing the XS Advanced Application Lifecycle Management Graphical User Interface

- To access the XS Advanced Application Lifecycle Management, choose one of the following options:
 - Use the following URL: `https://<server>:53280/index.html`
53280 is the default port for the XS Advanced Application Lifecycle Management graphical user interface when port-based routing is used.
If the routing configuration was changed to hostname routing during the installation, the URL can look different. For more information on routing configuration, see the *SAP HANA Server Installation and Update Guide* and SAP Note [2245631](#). In this case, you can check for the URL to access XS Advanced Application Lifecycle Management in the following places:
 - Using the XS advanced command line interface (CLI):
Use the `xs version (xs -v)` command. Below the information about the installed server version, under *Registered Service URLs*, it shows the *product-installer* URL.
 - Using the SAP HANA XS Advanced Cockpit:
In the tile catalog, choose the *Application Monitor* tile. In the list of applications (opening may take some time), locate the *product-installer-ui* and choose the *URL* link.

- In SAP HANA Cockpit, for a specific resource, choose one of the links available in the *Application Lifecycle Management* section:
 - For software components: *Install, update and uninstall XS advanced components*
 - For products: *Install, update and uninstall XS advanced products*
 - To display installed software components and products: *Show history*
 - In the XS Advanced Cockpit, from the *Spaces* in the *Organizations* section, choose the *SAP* space. In the *SAP* space, locate the link to the *product-installer-ui* and start the UI by clicking on the link displayed under *Application Routes*.
The *SAP* space is only visible if the user is assigned to it as Space Manager.
2. Log in with your user credentials.
For more information, see [Prerequisites and Authorizations \[page 961\]](#).
 3. Select the space in which you would like to work by choosing *Switch Space* under *Additional Functionality*.

Note

When you use the XS Advanced Application Lifecycle Management graphical user interface for the first time, you are asked to select a space in which you would like to work. Select one from the list of spaces or use the search functionality on top of the list of spaces. This list displays only those spaces that you are authorized for. You can change the space that you are working in later on at any time.

Elements of the XS Advanced Application Lifecycle Management Graphical User Interface

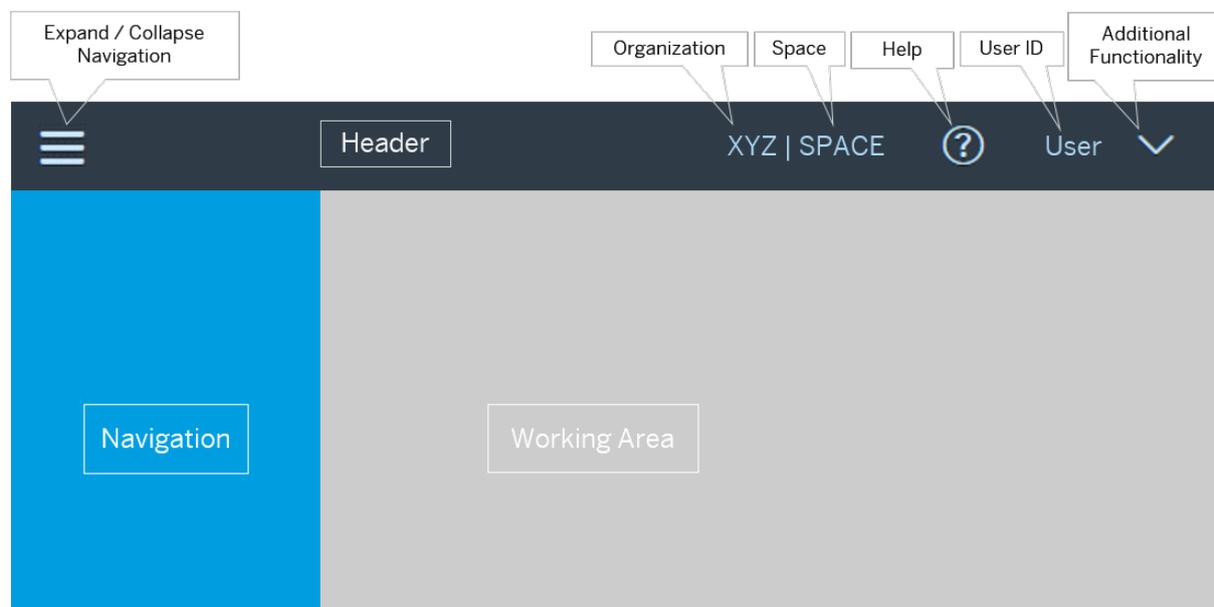
In the header of the graphical user interface on the right, you can find the name of the space and the SID of the system that you are currently working with. The questionmark leads you to detailed information about how to use the XS Advanced Application Lifecycle Management graphical user interface. By clicking on your user ID, additional functions like switching spaces and logging out are offered.

On the left side, you can decide whether you would like to work on software components or products or would like to take a look into the installation history. In the working area, you can see details for the topic that you chose:

- **Software Components**
A list of the software components that are already installed in the selected space is shown. There are several options to manage the installed software components. If no software components are installed in the selected space, the list is empty.
- **Products**
On the left, a list of products that are already installed in the selected space is shown. On the right, details for the product selected in the list are provided. If no products are installed in the selected space, the list is empty.
- **Installation History**
The installation history provides an overview of all installation, update and uninstallation activities for both software components and products in the selected space. It comprises detailed information about the activities, for example, type and status of the activity, name of the product or software component, version before and after the activity was executed. For a detailed installation, update or uninstallation screen with all actions that were executed, you can click on an activity. You can also download the log if you need it for further analysis, go to the product concerned or to the screen of software components to see the current state.

To find activities for a certain user or to search for a process ID, you can use the search above the list. The menu to the right of the search field offers different sorting options.

The following figure shows the elements of the XS Advanced Application Lifecycle Management graphical user interface:



Elements of the XS Advanced Application Lifecycle Management Graphical User Interface

Related Information

[The Multi-Target Application Model](#)

8.2.2.3.1 Install SAP HANA Software Components

You can use the XS Advanced Application Lifecycle Management graphical user interface to easily install new software components on your SAP HANA system.

Prerequisites

- The required installation file is available at a file location that you can reach from your computer.
- Make sure that you have selected the correct space.

Procedure

1. In the navigation area on the left, choose *Software Components* and then *Install/Update* in the working area.

2. Browse for the installation archive and, if required, enter one or more extension descriptors.
3. To upload the selected file, choose *Continue*.

i Note

If you want to clear the file selection, choose *Reset*. If you want to cancel the installation process, choose *Cancel*.

A list of the uploaded software components with some detailed information (for example, currently installed version, uploaded version, status) is displayed. The status for the software component that you uploaded is *Installation*.

4. Start the installation.

In the working area, you can now find information about the software component, its version, installation time and so on. You can also follow the installation progress which is done in the three steps *Validation*, *Deployment* and *Registration*. The icons for the different steps will change after a step is completed. You can click on the respective step for detailed information.

As soon as a green hook is shown for each step, the installation is successfully done.

Results

The new software component is successfully installed on your SAP HANA system. You can check this by refreshing the list of installed software components.

8.2.2.3.2 Update SAP HANA Software Components

You can use the XS Advanced Application Lifecycle Management graphical user interface to easily update software components that are already installed on your SAP HANA system.

Prerequisites

- Make sure that the installation file is available at a file location that you can reach from your computer.
- Make sure that you have selected the correct space.
- A previous version of the software component has already been installed to the selected space on your SAP HANA system.
- The software component is not part of a product.

i Note

A software component can only be updated if it is not part of a product. If it is part of a product, you need to update the complete product.

Procedure

1. In the navigation area on the left, choose *Software Components*.

A list of the software components that are already installed in the selected space is shown.

i Note

If the list is empty, there are no software components installed in the selected space and you cannot update any software component. Make sure that you selected the correct space.

2. In the working area, choose *Install/Update*.
3. Browse for the installation archive and, if required, enter an extension descriptor.
4. To upload the selected file, choose *Continue*.

i Note

If you want to clear the file selection, choose *Reset*. If you want to cancel the update process, choose *Cancel*.

A list of the uploaded software components with detailed information (for example, currently installed version, uploaded version, status) is displayed. Depending on the installed and uploaded versions of the software component, the following situations are possible:

- If the version of the software component that you uploaded is higher than the one that is already installed, the status displayed in the list is *Update*. You can start the update directly.
 - If the version that is already installed is the same as the one that you uploaded, you cannot start the installation directly. The status for the selected software component is *Already installed, to overwrite set the corresponding option*. To do so, choose *Options* in the upper right side of the working area and select *Overwrite the same version of software component*. This can be helpful if you would like to repair the installed version. In the list of uploaded software components, the status is changed to *Overwrite*. You can now start the update process.
 - If the version of the software component that is already installed is higher than the one that you just uploaded, you cannot proceed. You cannot downgrade software components during the update process.
5. To start the update process, choose *Start Installation*.

In the working area, you can now find information about the software component, its version, installation time and so on. You can also follow the installation progress which is done in the three steps *Validation*, *Deployment* and *Registration*. The icons for the different steps will change after a step is completed. You can click on the respective step for detailed information.

As soon as a green hook is shown for each step, the installation is successfully done.

Results

The new version of the software component is successfully installed on your SAP HANA system. You can check this by refreshing the list of installed software components.

8.2.2.3.3 Install SAP HANA Products

You can use the XS Advanced Application Lifecycle Management graphical user interface to easily install new products on your SAP HANA system.

Prerequisites

- The required installation file is available at a file location that you can reach from your computer.
- Make sure that you have selected the correct space.

Procedure

1. In the navigation area on the left, choose *Products* and then *Install/Update* below the list of installed products in the working area.
2. Browse for the installation archive and, if required, enter an extension descriptor.
3. To upload the selected file, choose *Continue*.

i Note

If you want to clear the file selection, choose *Reset*. If you want to cancel the installation process, choose *Cancel*.

Details about the uploaded product installation file (for example, vendor, version) are displayed. In addition, you can find information about the product instances that are part of the installation file. You can either decide to install the complete product with all instances or you can select one or several instances for installation.

i Note

It might happen that software components that are part of the product to be installed are already installed on your system. In this case you have to decide on how to proceed with these software components before you can start the installation. To do so, choose *Options* in the upper right side of the working area and maintain the settings.

4. Start the installation.

In the working area, you can now find information about the product, its version, installation time and so on. You can also follow the installation progress which is done in the three steps *Validation*, *Deployment* and *Registration*. The icons for the different steps will change after a step is completed. You can click on the respective step for detailed information.

As soon as a green hook is shown for each step, the installation is successfully done. You can now download the installation log.

Results

The new product is successfully installed on your SAP HANA system. You can check this by refreshing the list of installed products.

8.2.2.3.4 Update SAP HANA Products

You can use the XS Advanced Application Lifecycle Management graphical user interface to easily update products that are already installed on your SAP HANA system.

Prerequisites

- Make sure that the installation file is available at a file location that you can reach from your computer.
- Make sure that you have selected the correct space.
- The product has already been installed to the selected space on your SAP HANA system.

Procedure

1. In the navigation area on the left, choose *Products*.

A list of the products that are already installed in the selected space is shown.

i Note

If the list is empty, there are no products installed in the selected space and you cannot update any product. Make sure that you selected the correct space.

2. In the working area below the list of installed products, choose *Install/Update*.
3. Browse for the installation archive and, if required, enter an extension descriptor.
4. To upload the selected file, choose *Continue*.

i Note

If you want to clear the file selection, choose *Reset*. If you want to cancel the update process, choose *Cancel*.

A preview shows detailed information (for example, vendor, currently installed version, uploaded version) about the uploaded product and a list of product instances that are part of the uploaded installation file. You can either decide to update the complete product with all instances or you can select one or several instances for the update. The overview of instances also shows if software components of the respective instance are already installed. By default, instances that are already installed are selected for update. You cannot deselect them. To find out more about the software components concerned, you can click on the name of the instance.

Depending on the installed and uploaded versions of the product and its software components, the following situations are possible:

- If the versions of the product and all its instances and software components that you uploaded are higher than the versions of those that are already installed, you can start the update directly.
 - If an instance or some software components are already installed with the same version that is part of the uploaded file, these software components are not reinstalled automatically. To install them, choose *Options* in the upper right side of the working area and select *Overwrite the same version of software component*. This can be helpful if you would like to repair installed versions. In the list of instances, the *Prerequisite Check* is changed to *Same version will be overwritten*.
 - If the version of the software component that is already installed is higher than the one that you just uploaded, you would downgrade the already installed software components which is not possible during the update process. In this case, choose *Options* in the upper right side of the working area and select *Keep newer version of software component*.
 - You can select additional instances of the product for installation during the update process.
5. To start the update process, choose *Start Installation*.

In the working area, you can now find information about the software component, its version, installation time etc. You can also follow the installation progress which is done in the three steps *Validation*, *Deployment* and *Registration*. The icons for the different steps will change after a step is completed. You can click on the respective step for detailed information.

As soon as a green hook is shown for each step, the installation is successfully done.

Results

The new version of the software component is successfully installed on your SAP HANA system. You can check this by refreshing the list of installed software components.

8.2.2.3.5 Uninstall SAP HANA Products or Software Components

You can use the XS Advanced Application Lifecycle Management graphical user interface to easily uninstall products or software components from your SAP HANA system.

Prerequisites

The software component that you want to delete must not be part of an installed product.

i Note

If the software component that you want to delete is part of an installed product, you cannot uninstall it unless you uninstall the complete product or the product instance that contains the software component.

Procedure

1. In the navigation area on the left, choose *Software Component* or *Product* according to your needs.

A list of installed software components or products is shown. If you want to uninstall an individual instance of a product, you can select it in the details of the selected product.

2. Choose *Uninstall* for the software component or product that you would like to uninstall.

You are prompted to confirm the uninstallation of the selected software component or product and to decide whether or not to delete services that were created during the previous deployment. These services might contain important user data, for example database content, that cannot be recovered after it is removed. If you choose to keep the services but want to delete them later, you can use the `xs delete-service` command in the command line client. For an overview of all existing services in a given space and their bound applications, use the `xs services` command.

3. To uninstall the software component or product and the related services if selected, choose *Uninstall*.

In the working area, you can now find information about the software component or product, its version, installation time and so on. You can also follow the uninstallation progress which is done in the three steps *Validation*, *Undeployment* and *Deregistration*. The icons for the different steps will change after a step is completed. You can click on the respective step for detailed information.

As soon as a green hook is shown for each step, the uninstallation is successfully done.

Results

The software component or product is successfully uninstalled from your SAP HANA system. You can check this by refreshing the list of installed software components or products.

8.2.3 Configuring SAP HANA Applications with the Process Engine

The Process Engine (PE) is a framework available with SAP HANA application lifecycle management to enable automated technical configuration.

After the installation of a product or a delivery unit, an application typically must be configured before it can be used. The configuration tasks are described in the installation guides that are provided on the SAP Help Portal (help.sap.com). Instead of performing cumbersome and error-prone manual activities, you can use the Process Engine to automate application configuration completely or partially. As a prerequisite, your application must provide content for the automated technical configuration.

The Process Engine (PE) framework is installed with SAP HANA application lifecycle management as automated content. It is available from the following locations:

- On the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/lm/pe/ui/`
- Using the *Configuration Services (Process Engine)* tile in SAP HANA Application Lifecycle Management XS user interface

- Using the *Manage Configuration Services* tile in SAP HANA cockpit

Related Information

[Tutorial: Execute a Configuration Service with Process Engine \[page 991\]](#)

[Process Engine Roles \[page 990\]](#)

[Troubleshooting \[page 993\]](#)

8.2.3.1 Process Engine Roles

To grant users the privileges they require to perform tasks with the Process Engine, you must assign them the relevant Process Engine roles.

The following table lists the roles that are available for tasks related to the Process Engine. The roles are hierarchical and interlinked. The `sap.hana.xs.lm.roles::Administrator` role is the *Administrator* role of SAP HANA application lifecycle management and grants the privileges of all other Process Engine-related roles as well as application lifecycle management roles. For more information, see *SAP HANA Application Lifecycle Management Roles* in the *SAP HANA Application Lifecycle Management Guide*.

→ Recommendation

As repository roles delivered with SAP HANA can change when a new version of the package is deployed, either do not use them directly but instead as a template for creating your own roles, or have a regular review process in place to verify that they still contain only privileges that are in line with your organization's security policy. Furthermore, if repository package privileges are granted by a role, we recommend that these privileges be restricted to your organization's packages rather than the complete repository. To do this, for each package privilege (`REPO.*`) that occurs in a role template and is granted on `.REPO_PACKAGE_ROOT`, check whether the privilege can and should be granted to a single package or a small number of specific packages rather than the full repository.

Roles available for the Process Engine

Role	Description
<code>sap.hana.xs.lm.pe.roles::PE_Display</code>	The user can monitor processes and display services.
<code>sap.hana.xs.lm.pe.roles::PE_Execute</code>	In addition to the previous role, the user can start, stop, skip, and resume processes.
<code>sap.hana.xs.lm.pe.roles::PE_Activate</code>	In addition to the previous roles, the user can activate services from repository files.
<code>sap.hana.xs.lm.roles::Administrator</code>	The user can install products. This role includes all previous roles.

8.2.3.2 Tutorial: Execute a Configuration Service with Process Engine

In this tutorial, you use the demo content delivered with the Process Engine to execute a configuration service.

Prerequisites

- An SAP HANA system is available.
- SAP HANA XS is up and running on the SAP HANA system.
- Depending on the task you want to perform with the Process Engine, you must have the privileges based on a role granted by one of the Process Engine role templates described in *Process Engine Roles*. The link to the topic is in the *Related Information* section. The privileges of the `sap.hana.xs.lm.pe.roles::PE_Activation` role allows you to perform all Process Engine tasks.

Context

The Process Engine uses different terms for identifying design time or runtime artifacts. The *service* is the core entity at design time. It has multiple attributes describing its purpose and steps representing the executable entities. They perform the actual work during execution. An executable can be a JavaScript function in an XS JavaScript library or an SQL stored procedure. When starting a service, the Process Engine creates a *process* based on a service. It copies all steps associated with the service as tasks, and it copies the parameters of the selected variant to the parameters of the process. Furthermore, the Process Engine associates a *status* with the process.

You execute the following steps to configure the demo service:

- Activate the demo service.
Services are delivered as repository objects. The services required by the administrator need to be enabled once before use. This activity is called *activation*.
- Prepare the demo service parameters.
The demo service needs parameters during execution. The set of required parameters is stored under a common key, the *variant*. Before you can start a service you need to prepare variants. Since you are about to start the service for the first time, you do not have any variants prepared. If you repeat an execution, you can use an existing variant. For the demo service, you enter *user* and *password*. Since this is a demo example, the user does not need to exist and the password can be any set of characters.
- Start the demo service.
The demo service consists of the following steps:
 - JS_APPVAR by JavaScript
This step executes a JavaScript function that shows how to consume and return parameters in JavaScript.
 - SQL_APPVAR by SQL Script
This step executes a SQL script function that shows how to consume and return parameters in SQL script.

i Note

The demo content does not perform any configuration of the system. It only writes messages into the log of the Process Engine. It provides you with a hands-on experience for using the Process Engine.

Procedure

1. Open SAP HANA Application Lifecycle Management.

SAP HANA Application Lifecycle Management is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/lm`

2. Choose *Configuration Services (Process Engine)*.

The process engine opens in a new browser window or a tab.

3. Select *Services* on the left-hand side of the screen.

A list of services available for configuration appears. Active services are indicated by a green status icon, inactive services have a grey status icon. Inactive services must be activated before they can be started.

You find the demo service *DEMO_VarCont* as inactive in the list.

4. Select *DEMO_VarCont* and choose *Activate*.

Note that the activation process can take some time. After the service was activated successfully, you can start it.

If the activation were not successful, you can find the error messages in a detailed log.

5. To display the service after the activation, choose *Go To Service*.

The details of the active service *Demo Service with Process Engine Variant Container* appear.

6. To prepare the parameters for the demo service, choose *Maintain Variants*.

A new screen for variant maintenance appears.

7. Enter a *user* name and a *password* as parameters and choose *Save As*.
8. Enter a *variant ID* and, optionally, a *description*, and then choose *Create*.
9. Return to the previous screen to view the variant that you just created.

If you open the *Steps* tab on screen, two executable steps are displayed.

10. To start the service, select your variant and choose *Start Variant*.

The *Process* tab opens and a new process appears at the top of the list.

11. Select the process to navigate to the process details.

A list of tasks appears.

12. Choose *Refresh* to observe the progress of the process.

The overall process status is displayed on top of the progress bar. It is a cumulation of the statuses of the individual tasks.

The status icons allow you to intervene in the process execution if errors occur. You can click on an icon to display the task log.

13. When the process completion reaches 100%, choose the *Parameters* tab.

You see an overview on the scalar parameters and their changes during execution.

Entries for the step *NA* show the parameter values after the variant container is copied and before the process execution starts. The other entries show the values after the step was executed.

14. To view the log of the *Consuming and returning parameters via SQL* task, select this task at the bottom of the screen.
 - a. Search for a message with a green status that starts with *JavaScript function sends*. At the end of the message, you see the parameter value of your **user**.
 - b. Search for a message with a green or orange status that contains the text ... *password received*.

The step compares the received value of the parameter you entered as **password** with a value set by the demo service. If you entered the password as set by the demo service, the Process Engine issues the message *Correct password received*. If you entered a different password, the Process Engine writes *Incorrect password received* in the log.

Results

You have used the demo configuration service of the Process Engine. You have activated the demo service, prepared the parameters, and executed the service. Afterward, you have checked the logs of the Process Engine.

Related Information

[Process Engine Roles \[page 990\]](#)

[Troubleshooting \[page 993\]](#)

8.2.3.3 Troubleshooting

If a process stops with errors, you should first analyze the logs to find out why an error occurred. Afterward, you have various options to respond to the error situation.

The Process Engine provides a process log and a task log. If a single task has an error you can start with the task log to analyze if an error message is due to a specific step. If this does not help, you can open the process log and search or filter for error messages.

- **Process Log**

This is a collection of all task logs and additional entries related to the process. This log contains all messages with technical details, including the log of the internal activities of the Process Engine. You can find this log when you open the *Log* tab in the single process view.

- **Task Log**

This is the log of the execution of a single task. Messages with technical details about the Process Engine usually are not displayed here. You get this log when you navigate to a task view by clicking on a task in the single process view.

You have the following options to respond to an error:

- If the error is only temporary or you solved the error already, you can execute the step again by choosing [Resume](#).
- You can decide to perform the task manually and skip the execution of the task by choosing [Skip](#).
- You can cancel the current process and start a new one. To do this, choose [Cancel](#).
- If you cannot resolve the error, and you need to contact SAP, open an incident and assign it to the support component of the application that provides the configuration content or, alternatively, to component HAN-LM-APP. Make sure that you attach the diagnosis information that you can download for each process using the link on the [Diagnosis Information](#) tab.

8.3 SAP HANA Content

SAP HANA content is structured in the way that delivery units (DUs) are used to group SAP HANA content artifacts (such as analytic, attribute or calculation views, and SQLScript procedures).

DUs are grouped to SAP HANA products in order to ship and install SAP HANA applications with all dependent artifacts (grouped in DUs). To distribute SAP HANA content, a product archive (*.ZIP file) or a delivery unit archive (*.tgz file) is used. There are various ways of acquiring and deploying these archive types.

SAP HANA content, which is developed on SAP HANA Extended Application Services (SAP HANA XS), classic model, can also be grouped in a DU.

For more information about SAP HANA content deployed automatically during platform installation or upgrade, see *Components Delivered as SAP HANA Content* in the *SAP HANA Security Guide*.

8.3.1 SAP HANA Archive Types

The difference between the various archive types is their method of deployment, and when the content is deployed.

The following archive types are available:

- **Product archive file (*.ZIP)**
A product version archive is a *.ZIP file containing 1-n software component archive files and the following metadata files: `stack.xml`, `pd.xml`. A software component archive file is created for each DU containing its archive file (*.tgz).
A product is usually the entity that delivers SAP HANA applications, but it can also be used for transports. SAP HANA content that can be downloaded independently is shipped as SAP HANA products in SAP HANA product archives. SAP HANA content that is not part of the SAP HANA database is called SAP HANA content add-on (or SAP HANA product). SAP HANA content add-ons are developed as part of the SAP HANA platform or as part of an application that runs on top of SAP HANA.
For information about how to deploy a product archive, see *Deploy a Product Archive (*.ZIP)*.
- **Software Component Archive (*.ZIP)**
A software component archive is a *.ZIP file (in previous versions also *.SAR files were delivered as software component archives) containing one delivery unit archive file (*.tgz) and (optionally) a

corresponding translation DU and the metadata file `SL_MANIFEST.XML`. A software component archive can be deployed with the same tool as product archives.

For information about how to deploy a software component archive, see *Deploy a Product Archive (*.ZIP)*.

- **Delivery unit archive file (*.tgz)**

A delivery unit archive is a *.tgz file containing the SAP HANA content artifacts that are created in the SAP HANA repository. A DU is used to deliver one or more software components from SAP (or a partner) to a customer.

For distribution using export/import and deployment, a DU is contained in a delivery unit archive (*.tgz file). It contains the objects and packages of a DU together with the metadata file `manifest.txt`. The transport is also offered at DU level.

The following types of delivery unit archive files are available:

- **Delivery unit archives as part of the SAP HANA database**

The following types of delivery unit archive files that are part of the SAP HANA database are available:

- **Automated content** is installed together with SAP HANA and imported into the SAP HANA repository during installation. This is an integral part of the SAP HANA database and is used by every SAP HANA database customer.

Automated content is located on the SAP HANA system in the following folder:

```
/usr/sap/<SID>/SYS/global/hdb/auto_content.
```

- **Non-automated content** is installed with SAP HANA, but needs to be imported into the SAP HANA repository manually by the system administrator. It is used for integral parts of the SAP HANA database, but is only used by a small number of customers.

Non-automated content is located on the SAP HANA system in the following folder:

```
/usr/sap/<SID>/SYS/global/hdb/content.
```

Delivery unit archives that are non-automated content of the SAP HANA database need to be deployed manually.

- **Independent delivery unit archives that are not part of the SAP HANA database**

Delivery unit archives that are not installed together with the SAP HANA database and are not part of the SAP HANA database need to be deployed manually.

For information about how to deploy or activate a delivery unit archive, see *Deploy a Delivery Unit Archive (*.tgz)*.

Related Information

[Deploy a Product Archive \(*.ZIP\) \[page 995\]](#)

[Deploy a Delivery Unit Archive \(*.tgz\) \[page 996\]](#)

8.3.2 Deploy a Product Archive (*.ZIP)

SAP HANA application lifecycle management provides a method of deploying a product archive file (*.ZIP file containing a product) or software component archive files (*.ZIP).

For more information, see *Installing and Updating SAP HANA Products and Software Components* in the *SAP HANA Application Lifecycle Management Guide*.

Related Information

[Installing and Updating SAP HANA Products and Software Components in SAP HANA XS Classic Model \[page 952\]](#)

8.3.3 Deploy a Delivery Unit Archive (*.tgz)

The following deployment methods for deploying a delivery unit archive file (*.tgz file containing a DU) are provided:

- SAP HANA Application Lifecycle Management
Choose ► [Products](#) ► [Delivery Units](#) ► [Import](#) .
This tool runs on the SAP HANA XS Web server.
For more information, see *Import a Delivery Unit* in the *SAP HANA Developer Guide (For SAP HANA Studio)*.
- SAP HANA Application Lifecycle Management
SAP HANA application lifecycle management provides functions for installing and updating SAP HANA products:
 - SAP Fiori application integrated in the SAP HANA Application Lifecycle Management XS application
 - `hdbalim` command line toolFor more information, see *Installing and Updating SAP HANA Products and Software Components* in the *SAP HANA Application Lifecycle Management Guide*.
- SAP HANA studio
Import function of the SAP HANA Modeler
Choose ► [File](#) ► [Import](#) ► [SAP HANA Content](#) ► [Delivery Unit](#) .

For more information, see *SAP HANA Modeling Guide*.

Related Information

[Installing and Updating SAP HANA Products and Software Components in SAP HANA XS Classic Model \[page 952\]](#)

9 Landscape Management and Network Administration

Manage your SAP HANA landscape and integrate SAP HANA into your network environment.

Related Information

[Landscape Management \[page 997\]](#)

[Network Administration \[page 1040\]](#)

9.1 Landscape Management

Manage your SAP HANA system landscape efficiently and respond flexibly to changing resource requirements.

Depending on your SAP HANA deployment model and landscape architecture, you can reconfigure and reorganize your system in a number of ways.

SAP HANA Tools

Copy and Move Operations

The following SAP HANA mechanisms allow you to copy and move systems and databases:

- **Platform LCM Tools:**
An SAP HANA system can be safely and efficiently reconfigured by decoupling the system hosts from the installation path through unregistration, and re-coupling them in a different configuration through registration. System reconfiguration tasks can be performed with the SAP HANA database lifecycle manager (HDBLCM).
- **SAP HANA System Replication:**
SAP HANA system replication can be used to create a copy of an SAP HANA system in a quick and simple way.
System replication mechanisms can also be used to copy and move tenant databases securely and conveniently from one SAP HANA system to another with near-zero downtime.
- **Backup and Recovery:**
You can create a homogeneous copy of an SAP HANA database by recovering an existing database to a different database. A homogenous database copy is a quick way to set up a cloned database, for example, for training, testing, or development.

System Rename

An SAP HANA system can be renamed by changing the system identifiers, like host names, SID, and instance number. Changing system identifiers can be performed with the SAP HANA database lifecycle manager (HDBLCM).

SAP Landscape Management

SAP Landscape Management is an add-on to SAP NetWeaver installed as an application with the SAP NetWeaver Application Server for Java (SAP NetWeaver AS for Java).

The enterprise edition of SAP Landscape Management software helps you to simplify and automate the efforts required to configure, provision, deploy, monitor, and manage your systems in both physical and virtualized infrastructures.

With SAP Landscape Management you can prepare, relocate, restart, start, stop, and unprepare single-host and multiple-host SAP HANA systems, and perform system replication operations.

Related Information

[Copying and Moving a System Using Platform LCM Tools \[page 999\]](#)

[Copying a System using System Replication \[page 1031\]](#)

[Copying and Moving Tenant Databases Between Systems \[page 1004\]](#)

[Copying a Database Using Backup and Recovery \[page 1374\]](#)

[Renaming a System \[page 1032\]](#)

[SAP Landscape Management, enterprise edition](#)

9.1.1 Copying and Moving a System Using Platform LCM Tools

An SAP HANA system can be safely and efficiently reconfigured by decoupling the system hosts from the installation path through unregistration, and re-coupling them in a different configuration through registration. System reconfiguration tasks can be performed with the SAP HANA database lifecycle manager (HDBLCM).

9.1.1.1 Relocate the SAP HANA System

It may become necessary to move the SAP HANA system to different hardware. If so, you need to unregister the SAP HANA system and re-register it on the new hardware. System relocation can be performed with the SAP HANA database lifecycle manager (HDBLCM).

Context

Relocation can be performed on both the entire SAP HANA system or on an individual SAP HANA instance. So, you can flexibly decide if you want to relocate only one host (for example, in the case of host outage) or relocate all hosts in the system (for example, in a system scale up).

i Note

An SAP HANA system can only be relocated to a target system that runs on the same hardware platform as the source system.

Procedure

1. Unregister the SAP HANA instance or the SAP HANA system.

- a. Log on to the SAP HANA source host.

If you are unregistering a SAP HANA multiple-host system, you can log on to any system host. If you are unregistering a multiple-host system and would like to unregister one instance at a time, perform the unregistration of each local host.

- b. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdb1cm
```

By default, <sapmnt> is /hana/shared.

- c. Start the unregister task:

- To unregister hosts using the SAP HANA database lifecycle manager command-line interface:
 - Start the command-line tool interactively:

```
./hdb1cm
```

and enter the index of the `unregister_instance` action, if you only want to unregister the local host from the SAP HANA system. Enter the index of the `unregister_system` action, if you want to unregister all hosts in the SAP HANA system. Or,

- To unregister hosts using the SAP HANA database lifecycle manager graphical user interface:
 1. Start the graphical user interface tool:

```
./hdblcmgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

2. Choose *Unregister SAP HANA System*.

- d. To continue with the task proceed as follows:

- In the command line interface: Enter *y*.
- In the graphical interface:
 1. To display the summary of the configuration data, choose *Next*.
 2. To execute the configuration task, choose *Run*. The system displays the configuration progress.
 3. After the configuration task has finished, you can:
 - View the log. To do so, choose *View Log*.
 - Exit the graphical user interface. To do so, choose *Finish*.

2. Mount the installation path (`sapmnt`), the datapath, and the logpath on the target hosts.

3. Register the new host.

- a. Log on to the SAP HANA target host.

If you are registering an SAP HANA multiple-host system, you can log on to any system host. If you are registering a multiple-host system and would like to register one instance at a time, perform the registration on the local host before the remote hosts.

- b. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblc
```

By default, `<sapmnt>` is `/hana/shared`.

- c. Start the register task:

- To register hosts using the SAP HANA database lifecycle manager command-line interface:
 - Start the command-line tool interactively:

```
./hdblc
```

and enter the index of the `register` and `rename` action, or

- To register hosts using the SAP HANA database lifecycle manager graphical user interface:
 1. Start the graphical user interface tool:

```
./hdblcmgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

2. Choose *Register and Rename SAP HANA System*.

- d. To continue with the task proceed as follows:

- In the command line interface: Enter *y*.
- In the graphical interface:
 1. To display the summary of the configuration data, choose *Next*.

2. To execute the configuration task, choose *Rename*. The system displays the configuration progress.
3. After the configuration task has finished, you can:
 - View the log. To do so, choose *View Log*.
 - Exit the graphical user interface. To do so, choose *Finish*.

Note

When using the command line, the options can be set interactively during configuration only if they are marked as interactive in the help description. All other options have to be specified in the command line. To call the help, in the SAP HANA resident HDBLCM directory of the SAP HANA system, execute the following command:

```
./hdbclm --action=unregister_instance --help
```

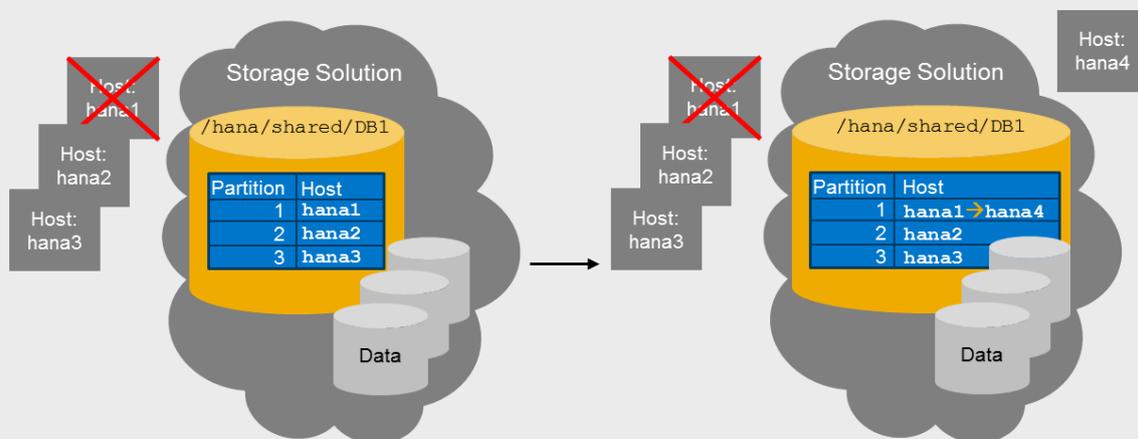
```
./hdbclm --action=unregister_system --help
```

```
./hdbclm --action=register_rename_system --help
```

Example

The following is an example of SAP HANA instance relocation from one host to another:

SAP HANA Instance Relocation



1. Unregister SAP HANA host hana1 using `--action=unregister_instance`

2. Mount `/hana/shared` on host hana4.
 3. Register SAP HANA host hana4 on the hana1 partition using `--action=register_rename_system` with host1 mapped to host4.

9.1.1.2 Copy or Clone an SAP HANA System

You can use the SAP HANA database lifecycle manager (HDBLCM) to make a copy or a clone of an SAP HANA system by copying the file system containing the SAP HANA database installation from an old storage solution to a new storage solution, and registering the copied SAP HANA system on new hosts.

Prerequisites

Before cloning the SAP HANA system, you must create a physical copy of the SAP HANA system (storage snapshot, file systems copy). The source system must be offline or a database snapshot must have been taken on the source system before the physical copy of the SAP HANA system is created.

i Note

An SAP HANA system can only be cloned or copied to a target system that runs on the same hardware platform as the source system.

Context

Cloning an SAP HANA system produces a new SAP HANA system, identical to the existing one. Copying an SAP HANA system produces a new SAP HANA system with the same landscape as the existing one, but slightly different parameter settings. If the interactive parameter defaults are accepted during host registration, the system is effectively cloned. If the new system parameters are set to different values, the new system is similar, but not identical to the source system.

You could, for example, copy an existing production system, and accept all parameter defaults during host registration except `system_usage`, which would be specified as "test". This configuration would allow you to have an almost identical copy of the existing system for test or quality assurance purposes.

⚠ Caution

Keep in mind that in a system copy refresh scenario all users and roles are overwritten in the target system.

Procedure

1. Copy the file system containing the SAP HANA database installation from the old storage solution to the new storage solution.
2. Mount the installation path (`sapmnt`), the data path, and the log path on the target hosts.
3. Register the new SAP HANA system on the target hosts.
 - a. Log on to the SAP HANA target host.

- b. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblc
```

By default, <sapmnt> is /hana/shared.

- c. Start the register task:

- To register hosts using the SAP HANA database lifecycle manager command-line interface:
 - Start the command-line tool interactively:

```
./hdblc
```

and enter the index of the `register` and `rename` action, or

- To register hosts using the SAP HANA database lifecycle manager graphical user interface:
 1. Start the graphical user interface tool:

```
./hdblcgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

2. Choose *Register and Rename SAP HANA System*.

- d. To continue with the task proceed as follows:

- In the command line interface: Enter *y*.
- In the graphical interface:
 1. To display the summary of the configuration data, choose *Next*.
 2. To execute the configuration task, choose *Run*. The system displays the configuration progress.
 3. After the configuration task has finished, you can:
 - View the log. To do so, choose *View Log*.
 - Exit the graphical user interface. To do so, choose *Finish*.

Note

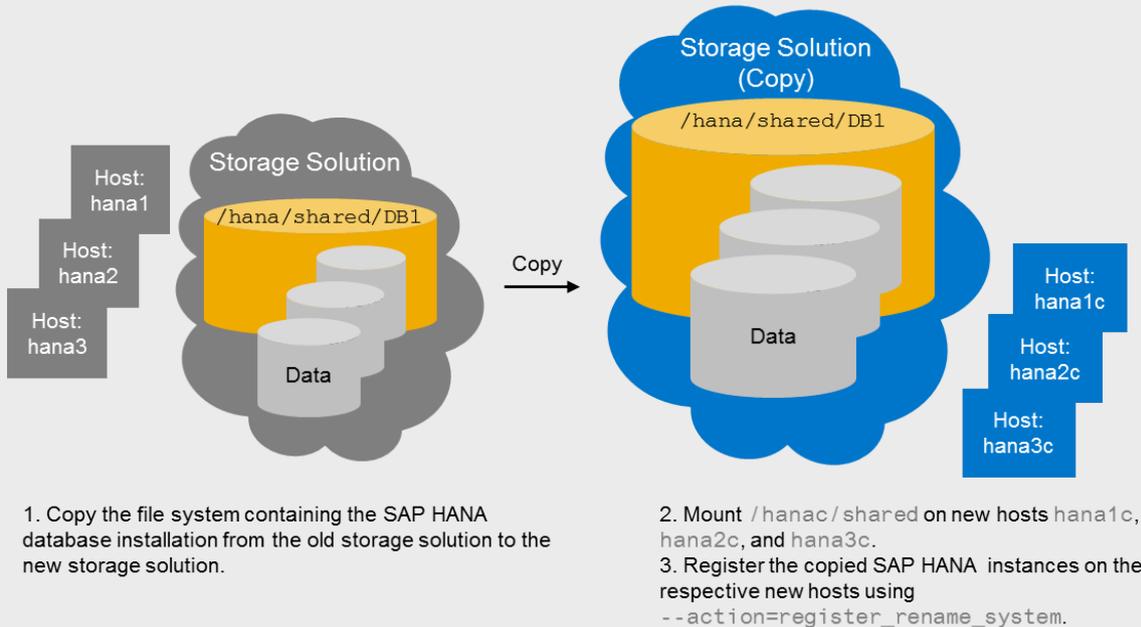
When using the command line, the options can be set interactively during configuration only if they are marked as interactive in the help description. All other options have to be specified in the command line. To call the help, in the SAP HANA resident HDBLCM directory of the SAP HANA system, execute the following command:

```
./hdblc --action=register_rename_system --help
```

❖ Example

The following is an example of an SAP HANA being cloned:

SAP HANA System Clone



9.1.2 Copying and Moving Tenant Databases Between Systems

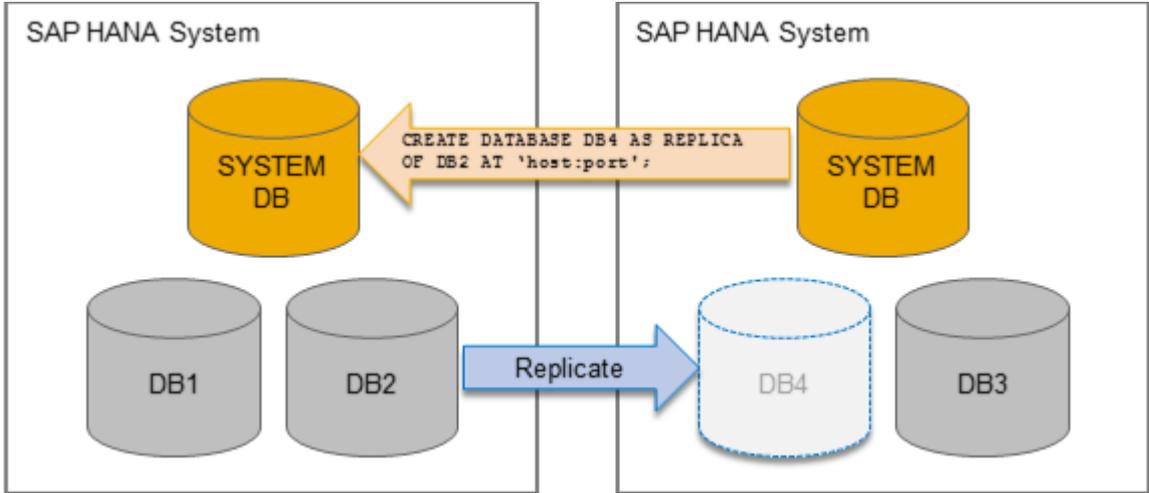
Using SAP HANA system replication mechanisms, SAP HANA tenant databases can be copied and moved securely and conveniently from one SAP HANA system to another with near-zero downtime. This allows you to respond flexibly to changing resource requirements and to manage your system landscape efficiently.

The following sections provide an overview of copying or moving a tenant database using system replication.

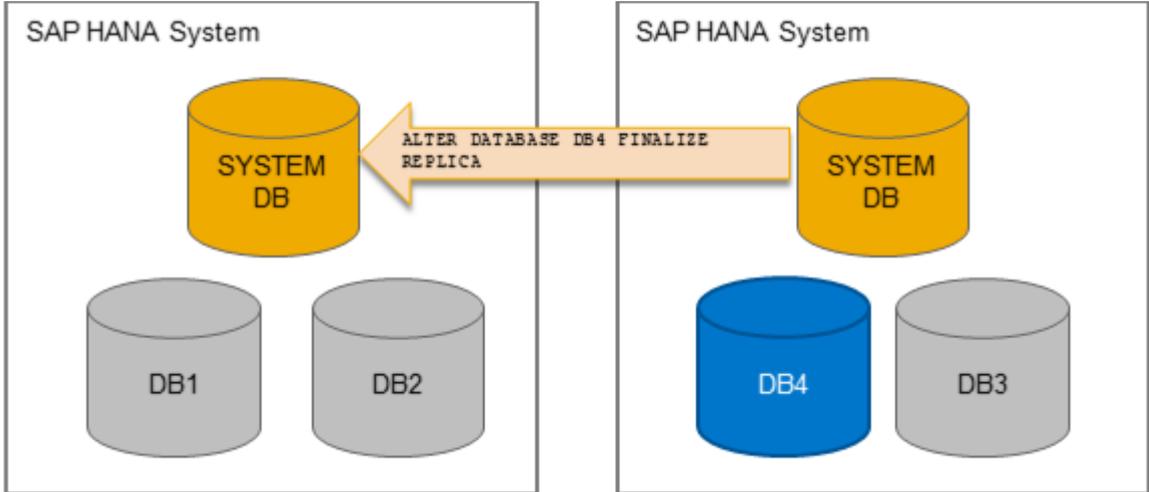
- [Process Overview \[page 1005\]](#)
- [Use Cases \[page 1006\]](#)
- [Which Data Is Copied or Moved? \[page 1006\]](#)
- [Recoverability After Copy or Move \[page 1007\]](#)
- [Client Communication After Move \[page 1008\]](#)
- [Prerequisites and Implementation Considerations \[page 1008\]](#)
- [Other Copy and Move Methods \[page 1008\]](#)

Process Overview

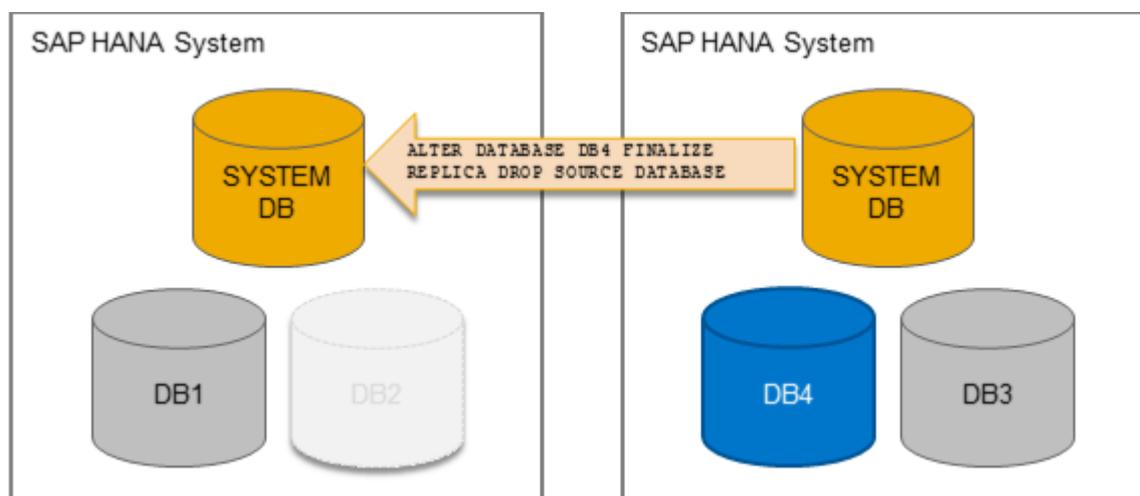
Copying and moving a tenant database are essentially the same process. First, the tenant database is copied through the replication of all of its data to a newly created tenant database in a target system:



Once all data has been successfully transferred, the new tenant database is started as a separate, independent database:



If the aim is to move the tenant database to the new system, the original tenant database is deleted and the new tenant database takes over:



The only difference between copying and moving a tenant database therefore is what happens to the original tenant database after all data has been transferred to the new tenant database in the target system.

In both cases, the new tenant database starts running as a fully separate, independent database.

Several tenant databases can be copied or moved to a system at the same time. It is also possible to copy or move a tenant database to a system with a different isolation level than the source system.

Use Cases

Copying and moving a tenant database from one system to another in this way has several applications, including:

- Load balancing between systems
For example, a tenant database is running a more demanding workload than anticipated, so you move it to a system running on a host with more CPU resources.
- Management of deployment environment
For example, you want to copy a tenant database running in your test system to the live production system.
- Tenant-database-specific upgrades
For example, you want to upgrade a single tenant database but not the entire system, so you move the tenant database to a system already running the higher version.
- Template databases
For example, you create a tenant database with a default configuration that you want to reuse as the basis for new tenant databases in other systems. You can simply copy the tenant database as a template to other systems.

Which Data Is Copied or Moved?

When a tenant database is copied or moved, data is replicated from the original tenant database to the new tenant database in the target system.

The following table indicates which types of data are replicated and which are not.

Type of Data	Replicated?
Data and logs of the tenant database	Yes
Trace and log files	No
Data backups	No
Configuration (*.ini) files with tenant-database-specific values	No
This refers to files in the directory <code>\$DIR_INSTANCE/.. /SYS/global/hdb/custom/config/<database_name></code>	
Certificates and certificate collections stored in the tenant database	Yes
This refers to the digital certificates and certificate stores used for certificate-based user authentication and secure communication between SAP HANA and JDBC/ODBC clients.	
<p>i Note</p> <p>If these certificates are stored in the file system in personal security environments (PSEs), they will not be replicated. To ensure that they are replicated, migrate the file-system-based PSEs to in-database certificate collections before copying or moving the tenant database. For more information about how to do this, see SAP Note 2175664.</p>	
Database-specific root key used for the internal data encryption service in the secure store file system (SSFS)	Yes
<p>i Note</p> <p>The root key used for data volume encryption is not replicated.</p>	
Database-specific root key used for backup encryption in the secure store file system (SSFS)	No (copy) / Yes (move)
Database-specific root key used for log encryption in the secure store file system (SSFS)	No (copy) / Yes (move)
Application function libraries	No

Recoverability After Copy or Move

When you copy a tenant database, the new tenant database does not have a backup history and cannot be recovered immediately after being copied. For this reason, it is important to perform a full data backup after you copy.

When you move a tenant database, the backup history of the original tenant database is retained in the new tenant database. As long as data and log backups of the source system are at a location accessible to the target system, the new tenant database is recoverable immediately after the move.

⚠ Caution

If you subsequently create a tenant database in the source system with the same name as the moved tenant database, the backup files of the original database are overwritten.

Client Communication After Move

We recommend that you do not specify physical host names in the SQL client connect string. Otherwise, you would have to reconfigure all of your applications after a move. Instead, configure virtual host names or virtual IP addresses. For more information on virtual IP addresses, see *Configure Host-Independent Tenant Addresses*.

Prerequisites and Implementation Considerations

- The copy and move process involves the creation of a new tenant database in the target system. Therefore, the target tenant database must not already exist in the target system.
- The target system must have a software version equal to or higher than the source system.
- If data volume encryption is enabled in the original system, data will be decrypted before replication and then re-encrypted (with a new root key) in the new database. However, during the copy and move process, data must be replicated via a secure (SSL/TLS) network connection by default.
- In a running system replication it is possible to copy or move tenant databases into a primary system or from a primary system into another target system different than the secondary system.
- There can be no changes to the topology of the original tenant database while the move or copy is in progress. In other words, until the copy or move has been finalized, it is not possible to add services to or remove services from the source tenant database.
- If the source system is configured for host auto-failover, the copy or process will fail in the event of failover to a standby host. If this happens, the new tenant database must be deleted on the target system and the copy or move process started again.
- The following components must not be configured in the source tenant database:
 - Rserve server
 - SAP HANA dynamic tiering (extended storage server)
 - SAP HANA accelerator for SAP ASE (extended transaction service)
 - SAP HANA streaming analytics (streaming host)

Other Copy and Move Methods

Backup and Recovery

It is possible to use backup and recovery to copy or move tenant databases between two systems. However, we recommend using SAP HANA system replication as described here. The main advantage of using system replication over backup and recovery is the absence of downtime. Using backup and recovery, you would have to shut down the original database after backing it up until the new database is successfully recovered. This is particularly critical if you are moving a tenant database. System replication is also a more convenient method because you don't need to move files between the different systems.

To copy or move a tenant database within the same system, we recommend using backup and recovery.

SAP HANA Database Lifecycle Manager (HDBLCM)

To copy or clone an entire system, use the SAP HANA database lifecycle manager (HDBLCM) as described in the *SAP HANA Platform Lifecycle Management* section of the *SAP HANA Administration Guide*.

Related Information

[SAP Note 2175664](#)

[Security of the Copy and Move Process \[page 1013\]](#)

[Copy and Move Process \[page 1009\]](#)

[Copy or Clone an SAP HANA System \[page 1002\]](#)

[Default Host Names and Virtual Host Names \[page 1069\]](#)

[Host Name Resolution for SQL Client Communication \[page 1074\]](#)

[Configure Host-Independent Tenant Addresses \[page 234\]](#)

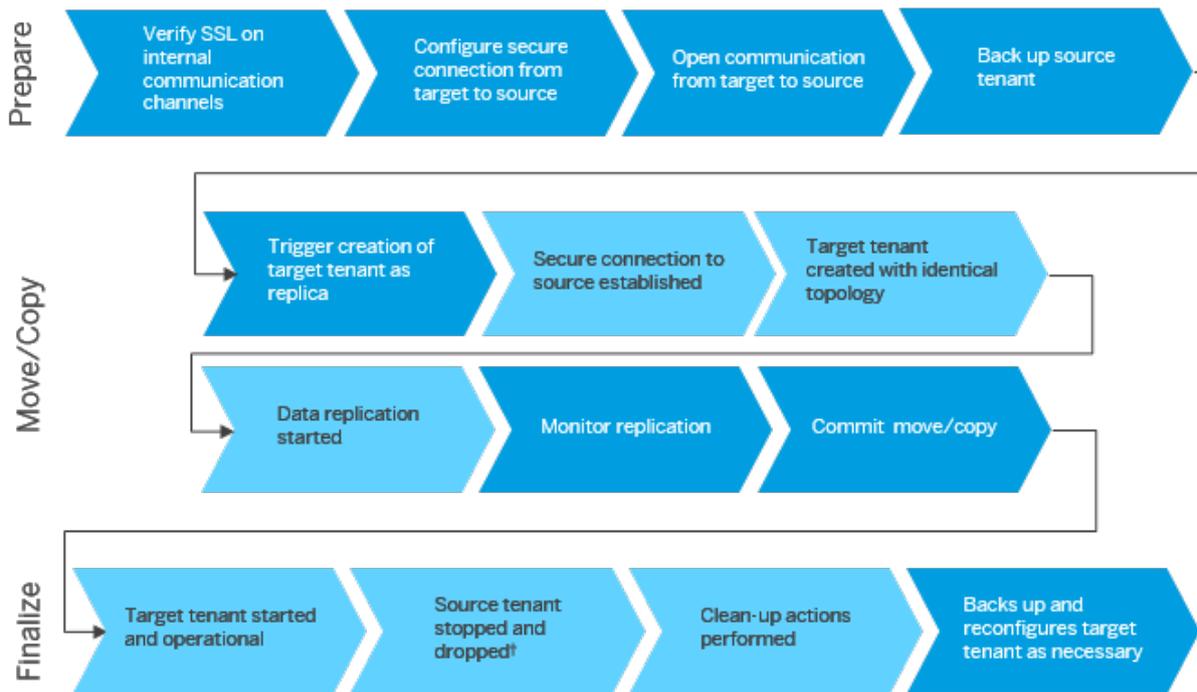
9.1.2.1 Copy and Move Process

Understand the stages and steps involved in copying and moving a tenant database from a source system to a target system using SAP HANA system replication.

Overview

The process of copying or moving a tenant database is driven entirely by the target system.

The following figure shows the stages involved, as well as who performs the individual steps in each stage: the system administrator or the target system. Each step is then described in more detail.



Legend

- Step performed by system administrator
- Step performed by system database of target system
- † Step performed only after a move

Tenant Move/Copy Process Flow

Prepare

Who?	Does What?	Where?
System administrator	Verifies that TLS/SSL is enabled on internal communication channels	System database of source and target system
	Configures secure connection from target system to source system by: <ol style="list-style-type: none"> 1. Creating a certificate collection with the purpose <code>DATABASE REPLICATION</code> and adding the root certificate of the source system to the new collection This will allow trust to be established between the system databases of the target and source system for external communication via SQL. 2. Creating a credential in the target system to enable authenticated access from the target system to the source system 	System database of target system

Who?	Does What?	Where?
	Opens communication from the target system to the source system by enabling source system services to listen on all network interfaces	System database of source system
	Creates a credential to enable authenticated access to the source system for the purpose of copying or moving a tenant database	System database of target system
	Backs up the source tenant database	System database of source system

For more information, see *Preparing to Copy or Move a Tenant Database*.

Copy and Move

Who?	Does What?	Where?
System administrator	Triggers the creation of the tenant database as a replica of the source tenant database by executing the SQL statement <code>CREATE DATABASE AS REPLICA</code>	System database of target system
System database of target system	Establishes a secure connection to the system database using the stored credentials created above For subsequent secure communication between the systems, a set of public and private key pairs and public-key certificates is generated in the source system database. These will be used to secure communication from the target system database to the source system database, and from the target tenant database to the source tenant database. The generated key pairs and certificates are imported into two newly created certificate collections in the target system database.	From target system to source system
	Creates a new tenant database with the same topology as the tenant database in the source system	Target system
	Initiates replication of data between the services in the source tenant database and the corresponding services in the target database	From source tenant database to target tenant database
System administrator	Monitors the progress of data replication in system view <code>SYS_DATABASES.M_DATABASE_REPLICAS</code>	System database of target system or source system

Who?	Does What?	Where?
	Once the replication status is <code>ACTIVE</code> (indicating that all data has been transferred), commits the copy by executing the SQL statement <code>ALTER DATABASE FINALIZE REPLICA</code>	System database of the target system
	In the case of a move, the administrator indicates that the source tenant database is to be dropped: <code>DROP SOURCE DATABASE.</code>	

For more information, see *Copy a Tenant Database to Another System* and *Move a Tenant Database to Another System*.

Finalize

Who?	Does What?	Where?
System database in target system	Starts the target tenant database	Target system
	If the source tenant database is being moved to the target system, stops and drops the source tenant database	Source system
	Performs clean-up operations: <ul style="list-style-type: none"> Deletes any cross-database dependencies to the original tenant database in other tenant databases of the source system (move only) Deletes any remote identity dependencies of users in the new tenant database in the target tenant database (copy and move) Generates a new root key used for data volume encryption and re-encrypts data if data volume encryption is enabled (copy and move) 	Source system and tenant database in target system
System administrator	Performs manual post-copy or post-move steps: <ul style="list-style-type: none"> Back up the target tenant database This is only necessary after a copy since the new tenant database does not have a backup history and cannot be recovered. After a move, the new tenant database has the backup history of the original tenant database and can be recovered if data and log backups of the source system are at a location accessible to the target system. Reverse preparatory steps required to secure the copy process If necessary, reconfigure parameters in *.ini files with tenant-database-specific values If necessary, reconfigure cross-database access 	System database of the target system

Cancel

Who?	Does What?	Where?
System administrator	Cancels the creation of the tenant database as a replica of the source tenant database by executing the SQL statement <code>DROP DATABASE <database_name></code>	System database of target system
System database of target system	Drops the target tenant database	Target system
	Performs clean-up operations	Target system

Related Information

[Security of the Copy and Move Process \[page 1013\]](#)

[Preparing to Copy or Move a Tenant Database \[page 1016\]](#)

[Copy a Tenant Database to Another System \[page 1024\]](#)

[Move a Tenant Database to Another System \[page 1027\]](#)

9.1.2.2 Security of the Copy and Move Process

Copying or moving a tenant database from one system to another is a secure end-to-end process.

Secure Network Communication

The copy and move process ensures end-to-end data encryption and host authentication on the basis of X.509 client certificates. Dedicated certificates and trust stores (referred to as certificate collections) are created as part of the copy or move process for the purpose of that specific copy or move. Certificates and certificate collections are stored directly in the system databases as database objects.

i Note

If secure network communication is not required, it can be disabled. As a result, a copy or move of a tenant database can be initiated without the exchange of certificates. For more information, see *Disable Secure Network Communication*.

For more information about in-database certificate management, see the *SAP HANA Security Guide*.

Authorization and Authentication

The copy and move process is triggered from the system database of the target system by a system administrator. To be able to execute the copy or move statements, the administrator user requires the system privilege `DATABASE ADMIN`.

To be able to establish a connection to the system database of the source system, the target system must be authenticated on the source system. This is achieved through the creation of a credential in the secure internal credential store of the system database of the target system. The required credential must be created manually by an administrator in the system database of the target system before the copy or move is started. For more information about the secure internal credential store, see the *SAP HANA Security Guide*.

Encryption Key Handling

SAP HANA features two data encryption services: data, backup, and log encryption in the persistence layer and an internal data encryption service available to applications requiring data encryption. The instance secure store in the file system (SSFS) is used to protect the root keys for these encryption services.

The root key used for data volume encryption is **changed automatically** in the new tenant database as part of the copy or move operation.

The root key used for the data encryption service is **not changed automatically** in the new tenant database as part of the copy or move operation. This root key is extracted from the SSFS of the original tenant database and replicated to the new tenant database and stored in the instance SSFS of the target system. It is not possible to change this root key manually.

⚠ Caution

Do not change the root key manually in the new tenant database. This will result in information in the SSFS and the database becoming inconsistent and encrypted data becoming inaccessible.

The root keys used for backup encryption and log encryption are **changed automatically** in the new tenant database as part of the copy operation.

The root keys used for backup encryption and log encryption are **not changed automatically** in the new tenant database as part of the move operation. These root keys are extracted from the SSFS of the original tenant database, replicated to the new tenant database and stored in the instance SSFS of the target system.

Cross-Database Dependencies

If cross-database access is enabled in the original tenant database, some configured dependencies are automatically deleted to ensure no unauthorized communication paths or user mappings can be exploited in the copied or moved tenant database.

Permitted Communication Paths

Part of the configuration of cross-database access involves specifying which tenant databases may communicate with each other and in which direction.

After a tenant database is moved to another system, it is deleted in the source system. However, communication paths referencing it will still exist in one or more of the other tenant databases in the source system. When the move operation is finalized, all such references to the original database are **automatically** deleted in other tenant databases of the source system.

Communication paths configured in the tenant database in the target system must manually be reconfigured after a move or copy.

User Mappings

Another aspect of cross-database access configuration is the mapping of users in one tenant database to users in another tenant database using remote identities.

After a tenant database is moved or copied to another system, some of its database users may still be associated as remote identities for users in other databases in the source system. When the move or copy operation is finalized, all remote identity information of users in the new tenant database is **automatically** deleted.

New user mappings must be manually configured in the tenant database in the target system.

Related Information

[Disable Secure Network Communication \[page 1015\]](#)

9.1.2.2.1 Disable Secure Network Communication

During the copy and move process, data is replicated via a secure (TLS/SSL) network connection by default. If this is not necessary in your scenario, you can disable this requirement.

Prerequisites

- You have the system privilege INIFILE ADMIN.
- Neither the source system or target system is configured for high isolation (the value of the `database_isolation` in the `[multidb]` section of the `global.ini` file is **low**).
- The value of parameter `ssl` in the `[communication]` section of the `global.ini` file is set to **off**.
- External communication is not configured for SSL in either the source system or target system.

Procedure

1. Open *Configuration of System Properties* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview*.

2. Use drop-down menus to select the *Configuration File* and the *Section* in order to display the [multidb] section of the `global.ini` file.
3. Select the edit icon for the parameter `enforce_ssl_database_replication` and change its value to **false**.
4. Enter the new value and click *Save*.

Related Information

[Modify a System Property in SAP HANA Cockpit \[page 299\]](#)

9.1.2.3 Preparing to Copy or Move a Tenant Database

Before you copy or move a tenant database to another system, you must perform several steps. These are primarily to enable the systems to communicate with each other securely.

Context

i Note

During the copy and move process, data is replicated via a secure (TLS/SSL) network connection by default. If you do not require a secure network connection and have disabled this feature, you can skip the first two steps: *Verify TLS/SSL Configuration of Internal Communication Channels* and *Set Up Trust Relationship Between Target and Source Systems*. For more information, see *Disable Secure Network Communication*.

1. [Verify TLS/SSL Configuration of Internal Communication Channels \[page 1017\]](#)
In both the source system and the target system, verify that TLS/SSL is enabled on internal communication channels on the basis of the system public key infrastructure (system PKI).
2. [Set Up Trust Relationship Between Target and Source Systems \[page 1018\]](#)
Create a certificate collection in the system database of the target system and add either the public-key certificate of the system database of source system, or the root certificate of the source system. This certificate is used to secure communication between the systems via external SQL connections.
3. [Open Communication From Target to Source System \[page 1021\]](#)
Open communication from the target system to the source system by enabling services in the source system to listen on all network interfaces.
4. [Create Credential for Authenticated Access to Source System \[page 1022\]](#)
Create a credential to enable authenticated access to the source system for the purpose of copying or moving a tenant database.
5. [Back Up Tenant Database \[page 1023\]](#)
Back up the tenant database that will be copied or moved.

Related Information

[Disable Secure Network Communication \[page 1015\]](#)

9.1.2.3.1 Verify TLS/SSL Configuration of Internal Communication Channels

In both the source system and the target system, verify that TLS/SSL is enabled on internal communication channels on the basis of the system public key infrastructure (system PKI).

Prerequisites

You have a user in the system database of both systems with the system privilege INIFILE ADMIN.

Context

During the copy and move process, data is replicated via a secure (TLS/SSL) network connection by default. If you do not require a secure network connection and have disabled this feature, you can skip this step. For more information, see *Disable Secure Network Communication*.

Procedure

→ Remember

This step must be performed in both the source system and the target system.

1. In the system database open *Configuration of System Properties* in SAP HANA cockpit by clicking the corresponding *Administration* link in the system *Overview*.
2. Use drop-down menus to select the *Configuration File* and the *Section* in order to display the [communication] section of the `global.ini` file.
3. Verify that value of the parameter `ssl` is set to **systemPKI** at the `SYSTEM` layer.
If it's not, change the value of the parameter accordingly.
4. Use drop-down menus to select the *Configuration File* and the *Section* in order to display the [system_replication_communication] section of the `global.ini` file.
5. Verify that value of the parameter `enable_ssl` is set to **on** at the `SYSTEM` layer.
If it's not, change the value of the parameter accordingly.

→ Tip

Alternatively, you can enable SSL on the basis of the system public key infrastructure by executing the following SQL statements:

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
 ('communication', 'ssl') = 'systemPKI' WITH RECONFIGURE;
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
 ('system_replication_communication', 'enable_ssl') = 'on' WITH
RECONFIGURE;
```

6. Restart the system.

Task overview: [Preparing to Copy or Move a Tenant Database \[page 1016\]](#)

Next task: [Set Up Trust Relationship Between Target and Source Systems \[page 1018\]](#)

Related Information

[Disable Secure Network Communication \[page 1015\]](#)

9.1.2.3.2 Set Up Trust Relationship Between Target and Source Systems

Create a certificate collection in the system database of the target system and add either the public-key certificate of the system database of source system, or the root certificate of the source system. This certificate is used to secure communication between the systems via external SQL connections.

Prerequisites

- You have a user in the system database of the target system with the system privileges CERTIFICATE ADMIN, TRUST ADMIN, and DATABASE ADMIN.
- You have a copy of the `extract_certificates.py` python script. The python script file must be accessible to the `<sid>adm` user. You can find the script attached to *SAP Note 2175664*.
- You have the public-key certificate of the system database of the source system (or the root certificate of the source system) used for external communication.

If this certificate does not already exist, you can create it using the SAPGENPSE tool or the SAP Web Dispatcher administration tool, both of which are delivered with SAP HANA. The certificate must be imported into the source system.

⚠ Caution

By default, SAP HANA allows encrypted communication for all exposed interfaces leveraging self-signed certificates. Although self-signed certificates allow communication encryption, full

communication security can only be reached leveraging certificates signed by a Certificate Authority (CA).

If the certificate does exist, its location depends on how you manage certificates in your system. Certificates stored in database (recommended) are contained in the certificate store. The required certificate is assigned to the collection with purpose `SSL`. Certificates stored in the file system are contained in tenant database-specific personal security environments or PSEs (default `$SECUDIR/sapsrv.pse`).

For more information, see *TLS/SSL Configuration on the SAP HANA Server* in the SAP HANA Security Guide and *Managing Client Certificates* in the SAP HANA Administration Guide.

Context

During the copy and move process, data is replicated via a secure (TLS/SSL) network connection by default. If you do not require a secure network connection and have disabled this feature, you can skip this step. For more information, see *Disable Secure Network Communication*.

Note

If you already have a CA-signed certificate, you can skip steps 1 through 3.

Procedure

1. Create a personal security environment (PSE) file using the `SAPGENPSE` tool.

```
sapgenpse gen_pse -p <path>/<file name>.pse -x "" -noreq "CN=<FQDN of source host>"
```

Example

```
sapgenpse gen_pse -p foo.pse -x "" -noreq "CN=sourcehost.domain"
```

2. Extract the generated private key and the self-signed certificate from the PSE file using the `extract_certificates.py` script.

```
python <path to script>/extract_certificates.py -p <file name>.pse
```

The script will print a list of one or more SQL statements that can be transferred to an SQL console using copy and paste.

3. In the system database of the source system, create a certificate collection and set its purpose to `SSL`. You can choose any name for the certificate collection.

You can do this using the *Certificate Collections* app of the SAP HANA cockpit or by executing the following SQL statements:

```
CREATE PSE <collection name>;
```

```
ALTER PSE <collection name> SET OWN CERTIFICATE '<private key and certificate>';  
SET PSE <collection name> PURPOSE SSL;
```

→ Tip

You can generate the ALTER PSE SQL statement using the `extract_certificates.py` script.

4. In the system database of the target system, create a certificate collection and set its purpose to DATABASE REPLICATION. You can choose any name for the certificate collection.

You can do this using the *Certificate Collections* app of the SAP HANA cockpit or by executing the following SQL statements:

```
CREATE PSE <collection name>;  
SET PSE <collection name> PURPOSE DATABASE REPLICATION;
```

5. If not already in the certificate store, import the public-key certificate of the system database of the source system (or the root certificate of the source system) into the certificate store of the target system.

You can do this using the *Certificate Store* app of the SAP HANA cockpit or by executing the following SQL statement:

```
CREATE CERTIFICATE FROM '<certificate content>';
```

6. Add the system database certificate (or root certificate) to the new collection.

You can do this using the *Certificate Collections* app of the SAP HANA cockpit or by executing the following SQL statement:

```
ALTER PSE <collection name> ADD CERTIFICATE <certificate id>;
```

→ Tip

You will find the certificate ID in the system view `SYS.CERTIFICATES`.

```
SELECT * FROM SYS.CERTIFICATES;
```

Task overview: [Preparing to Copy or Move a Tenant Database \[page 1016\]](#)

Previous task: [Verify TLS/SSL Configuration of Internal Communication Channels \[page 1017\]](#)

Next task: [Open Communication From Target to Source System \[page 1021\]](#)

Related Information

[SAP Note 2175664 - Migration of file system based X.509 certificate stores to in-database certificate stores](#) 
[Managing Client Certificates \[page 900\]](#)
[Disable Secure Network Communication \[page 1015\]](#)

9.1.2.3.3 Open Communication From Target to Source System

Open communication from the target system to the source system by enabling services in the source system to listen on all network interfaces.

Prerequisites

You have the credentials of operating system administrator `<sid>adm` for the source system.

Context

Use the SAP HANA database lifecycle manager (HDBLCM) to configure inter-service communication so that the services of the target system can listen on all available network interfaces.

i Note

It is only necessary to perform this step in the source system. However, if you later want to be able to monitor the progress of the copy or move operation from the source system, you can also do it in the target system.

Procedure

i Note

The following procedure describes how to do this using the Web user interface. For more information about using the command-line interface or graphical user interface of the SAP HANA database lifecycle manager, see the *SAP HANA Administration Guide*.

Instead of using the SAP HANA database lifecycle manager (HDBLCM), you can execute the following SQL statement:

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET ( 'communication',  
'listeninterface') = '.global' WITH RECONFIGURE;
```

1. Open the SAP HANA database lifecycle manager by entering the following URL in a browser:
`https://<host>:1129/lmsl/HDBLCM/<sid>/index.html`
2. Click the tile *Configure Inter-Service Communication*.
3. When prompted, enter the password of the `<sid>adm` user.
4. Select the setting *global*.
5. Click *Run* to apply the new setting.

6. Close the application and log out.

Task overview: [Preparing to Copy or Move a Tenant Database \[page 1016\]](#)

Previous task: [Set Up Trust Relationship Between Target and Source Systems \[page 1018\]](#)

Next task: [Create Credential for Authenticated Access to Source System \[page 1022\]](#)

Related Information

[Internal Host Name Resolution \[page 1072\]](#)

9.1.2.3.4 Create Credential for Authenticated Access to Source System

Create a credential to enable authenticated access to the source system for the purpose of copying or moving a tenant database.

Prerequisites

- You have a user in the system database of the target system with the system privilege CREDENTIAL ADMIN.

Context

You create a credential in the secure internal credential store of the system database of target system.

The credential store is used in SAP HANA to securely store credentials required for outbound connections. For more information about the secure internal credential store, see the *SAP HANA Security Guide*.

Procedure

In the system database of the target system, create a credential by executing the following SQL statement:

```
CREATE CREDENTIAL FOR COMPONENT 'DATABASE_REPLICATION' PURPOSE  
'<host:internal_port_of_system_DB_of_source_system>'
```

```
TYPE 'PASSWORD' USING
'user="<user_in_system_DB_of_source_system_with_DATABASE_ADMIN>";password="<password>"'
```

The values required for each parameter are as follows:

Parameter	Required Value
COMPONENT	DATABASE_REPLICATION
PURPOSE	Host name and internal port number of the system database of the source system
TYPE	PASSWORD
USING	User name of a user in the tenant database of the source system with the system privilege DATABASE ADMIN

Sample Code

```
CREATE CREDENTIAL FOR COMPONENT 'DATABASE_REPLICATION' PURPOSE
'host123456.acme.corp:30001' TYPE 'PASSWORD' USING
'user="DATABASE_ADMINISTRATOR";password="<password>"'
```

Task overview: [Preparing to Copy or Move a Tenant Database \[page 1016\]](#)

Previous task: [Open Communication From Target to Source System \[page 1021\]](#)

Next task: [Back Up Tenant Database \[page 1023\]](#)

9.1.2.3.5 Back Up Tenant Database

Back up the tenant database that will be copied or moved.

Context

You can back up the tenant database from the system database or from the tenant database directly. For more information, see *Creating Backups* in the *SAP HANA Administration Guide*.

Task overview: [Preparing to Copy or Move a Tenant Database \[page 1016\]](#)

Previous task: [Create Credential for Authenticated Access to Source System \[page 1022\]](#)

Related Information

[Creating Backups \[page 1313\]](#)

9.1.2.4 Copy a Tenant Database to Another System

Copy a tenant database from one SAP HANA system to another. The new copied tenant database runs as a separate, independent database.

Prerequisites

- All general system prerequisites are fulfilled. For more information, see *Copying and Moving Tenant Databases Between Systems*.
- All preparatory steps have been completed. For more information, see *Preparing to Copy or Move a Tenant Database*.
- You have a user in the system database of the target system with the system privilege `DATABASE ADMIN` and `CATALOG READ`.

Procedure

1. Create a tenant database in the target system as a copy of the original tenant database in the source system.

You do this by executing the `CREATE DATABASE` statement:

Code Syntax

```
CREATE DATABASE <target_database_name> [ AT [ LOCATION ]
'<target_hostname>[:<port_number_master_indexserver> ] ' ]
{ ADD '<servicetype>' [ AT [ LOCATION ]
'<target_hostname>[:<port_number_service> ]@<source_hostname>:<port_number_
service>' ] }...
{ AS REPLICA OF [ <source_database_name> ] AT [ LOCATION ]
'<source_hostname>[:<port_number_systemdb> ]' }
[ OS USER '<username>' OS GROUP '<groupname>' ]
[ NO START ]
[ <restart_mode> RESTART ]
```

Note

- As the location of the source tenant database, you specify the host name and port number for internal communication of the **system database** of the source system.
- If you enabled SSL, the host name must match the common name (CN) specified in the public-key certificate of the system database of source system.

- If you specify a service list, the number and type of services must match the source database.
- If your systems are configured for high isolation, you specify a valid OS user or OS group of the tenant database.
- Fallback snapshots are not copied to the target system.

Sample Code

```
CREATE DATABASE TARGET_DATABASE AS REPLICA OF SOURCE_DATABASE AT
'host123456.acme.corp:30001';
```

With the execution of this statement, the system database of the target system does the following:

- Establishes a secure connection to the system database of the source system
 - Creates a new tenant database with the same topology as the tenant database in the source system
 - Starts replicating data between the services in the source tenant database and the corresponding services in the target database
2. Monitor replication progress of data replication from the original tenant database to the new tenant database.

Use the system view `SYS_DATABASES.M_DATABASE_REPLICAS` to monitor the status of data replication in the system database of the target system or `SYS.M_DATABASES` to monitor directly in the new tenant database.

The current status of replication is shown in the field `REPLICATION_STATUS`. The value is aggregated across all individual services of the system, e.g. the system global status is only `ACTIVE`, if all individual services have replication status `ACTIVE`.

The following replication statuses are possible:

Status	Description
UNKNOWN	The secondary system did not connect to the primary system since the last restart of the primary system.
INITIALIZING	Data transfer is initialized. In this state, the secondary system cannot be used.
SYNCING	The secondary system is syncing again (e.g. after a temporary connection loss or restart of the secondary system).
ACTIVE	Initialization or sync with the primary system is complete and the secondary system is continuously replicating. If a crash occurs, no data will be lost in <code>SYNC</code> mode.
ERROR	A connection error occurred (details can be found in <code>REPLICATION_STATUS_DETAILS</code>).

The view `SYS_DATABASES.M_DATABASE_REPLICA_STATISTICS` provides detailed information about the replication process at the service level.

→ Tip

If the replication status is `ERROR`, use system view `SYS_DATABASES.M_DATABASE_REPLICA_STATISTICS` to investigate further.

If you cannot create a new replica because the source database is still in status "REPLICATING" even though the target database is already dropped, execute the statement `ALTER DATABASE`

```
<database_name> CANCEL REPLICA on the source system to clean up the system before re-attempting replication.
```

i Note

You can also monitor from the source system if you opened communication between the systems in both directions. For more information, see *Open Communication From Target to Source System*.

3. When replication status is `ACTIVE` (indicating that the new tenant database is in sync with the original tenant database), stop replication and finalize the copy by executing the following statement in the system database of the target system:

```
ALTER DATABASE <new_database_name> FINALIZE REPLICA
```

With the execution of the above statement, the system database of the target system performs the following actions in the new tenant database:

- Starts the new tenant database
- Changes the root key for data volume encryption and re-encrypts data in the new database if data volume encryption is enabled
- Deletes remote identities of database users if the original tenant database was configured for cross-database access

Next Steps

Perform the required manual post-move tasks.

Related Information

[Copying and Moving Tenant Databases Between Systems \[page 1004\]](#)

[Preparing to Copy or Move a Tenant Database \[page 1016\]](#)

[Perform Manual Post-Copy/Move Tasks \[page 1030\]](#)

9.1.2.5 Move a Tenant Database to Another System

Move a tenant database in one SAP HANA system to another. After a move, the original tenant database is deleted and the new tenant database takes over.

Prerequisites

- All general system prerequisites are fulfilled. For more information, see *Copying and Moving Tenant Databases Between Systems*.
- All preparatory steps have been completed. For more information, see *Preparing to Copy or Move a Tenant Database*.
- You have a user in the system database of the target system with the system privilege `DATABASE ADMIN` and `CATALOG READ`.

Procedure

1. Create a tenant database in the target system as a copy of the original tenant database in the source system.

You do this by executing the `CREATE DATABASE` statement:

Code Syntax

```
CREATE DATABASE <target_database_name> [ AT [ LOCATION ]
'<target_hostname>[:<port_number_master_indexserver> ] ' ]
{ ADD '<servicetype>' [ AT [ LOCATION ]
'<target_hostname>[:<port_number_service> ]@<source_hostname>:<port_number_
service>' ] }...
{ AS REPLICA OF [ <source_database_name> ] AT [ LOCATION ]
'<source_hostname>[:<port_number_systemdb> ]' }
[ OS USER '<username>' OS GROUP '<groupname>' ]
[ NO START ]
[ <restart_mode> RESTART ]
```

Note

- As the location of the source tenant database, you specify the host name and port number for internal communication of the **system database** of the source system.
- If you enabled SSL, the host name must match the common name (CN) specified in the public-key certificate of the system database of source system.
- If you specify a service list, the number and type of services must match the source database.
- If your systems are configured for high isolation, you specify a valid OS user or OS group of the tenant database.
- Fallback snapshots are not moved to the target system.

Sample Code

```
CREATE DATABASE TARGET_DATABASE AS REPLICA OF SOURCE_DATABASE AT  
'host123456.acme.corp:30001';
```

With the execution of this statement, the system database of the target system does the following:

- Establishes a secure connection to the system database of the source system
- Creates a new tenant database with the same topology as the tenant database in the source system
- Starts replicating data between the services in the source tenant database and the corresponding services in the target database

2. Monitor replication progress of data replication from the original tenant database to the new tenant database.

Use the system view `SYS_DATABASES.M_DATABASE_REPLICAS` to monitor the status of data replication in the system database of the target system or `SYS.M_DATABASES` to monitor directly in the new tenant database.

The current status of replication is shown in the field `REPLICATION_STATUS`. The value is aggregated across all individual services of the system, e.g. the system global status is only `ACTIVE`, if all individual services have replication status `ACTIVE`.

The following replication statuses are possible:

Status	Description
UNKNOWN	The secondary system did not connect to the primary system since the last restart of the primary system.
INITIALIZING	Data transfer is initialized. In this state, the secondary system cannot be used.
SYNCING	The secondary system is syncing again (e.g. after a temporary connection loss or restart of the secondary system).
ACTIVE	Initialization or sync with the primary system is complete and the secondary system is continuously replicating. If a crash occurs, no data will be lost in <code>SYNC</code> mode.
ERROR	A connection error occurred (details can be found in <code>REPLICATION_STATUS_DETAILS</code>).

The view `SYS_DATABASES.M_DATABASE_REPLICA_STATISTICS` provides detailed information about the replication process at the service level.

→ Tip

If the replication status is `ERROR`, use system view `SYS_DATABASES.M_DATABASE_REPLICA_STATISTICS` to investigate further.

If you cannot create a new replica because the source database is still in status "REPLICATING" even though the target database is already dropped, execute the statement `ALTER DATABASE <database_name> CANCEL REPLICA` on the source system to clean up the system before re-attempting replication.

i Note

You can also monitor from the source system if you opened communication between the systems in both directions. For more information, see *Open Communication From Target to Source System*.

- When replication status is `ACTIVE` (indicating that the new tenant database is in sync with the original tenant database), stop replication and finalize the move by executing the following statement in the system database of the target system:

```
ALTER DATABASE <new_database_name> FINALIZE REPLICA DROP SOURCE DATABASE
```

With the execution of the above statement, the system database of the target system performs the following actions:

- Starts the new tenant database
- Changes the root key for data volume encryption and re-encrypts data in the new database if data volume encryption is enabled
- Drops the original tenant database in the source system

i Note

To ensure that the new tenant database can be recovered to the most recent consistent state after the move, data backups are not deleted as part of the move process. This is important in the event that a backup is created in the original tenant database after replication has finished but before the original database is finally deleted.

- Deletes any communication paths configured for cross-database access that reference the original tenant database in the other tenant databases of the source system

Next Steps

Perform the required manual post-move tasks.

Related Information

[Copying and Moving Tenant Databases Between Systems \[page 1004\]](#)

[Preparing to Copy or Move a Tenant Database \[page 1016\]](#)

[Perform Manual Post-Copy/Move Tasks \[page 1030\]](#)

9.1.2.6 Perform Manual Post-Copy/Move Tasks

After you have committed the copy or move and the new tenant database is up and running, you must perform several manual tasks.

Procedure

1. Back up the new root keys to a root key backup file (*.rkb) in a secure location.

⚠ Caution

Store the root key backup file in a safe location. Losing this file may result in the database being unrecoverable.

2. Perform a full data backup of the new tenant database (copy only).
3. Reverse the preparatory steps required to secure the copy or move process:
 - Close network communication from the target system to the source system. See *Open Communication From Target to Source System*.
 - Delete the credential used by the system database of the target system to access the source system. See *Create Credential for Authenticated Access to Source System*.
 - Delete the certificate collection with purpose DATA REPLICATION created to secure external communication between the systems. See *Set Up Trust Relationship Between Target and Source Systems*.
 - If you created and signed the certificate yourself, delete the PSE file from the file system.
4. If necessary, reconfigure parameters in *.ini files with tenant-database-specific values. See *Configuration Parameters in Tenant Database Systems*.
5. Reconfigure cross-database access, if required. See *Enable and Configure Cross-Database Access*.
6. After configuring cross-database access, rebuild or repair cross-database objects. For improved performance, SAP HANA stores object IDs of remote objects in the catalog persistence. If a tenant database is copied or moved, these remote object IDs are likely to change and must be rebuilt by executing the CHECK_CATALOG procedure. The following objects are supported: functions, procedures, synonyms, SQL views, and calculation views.

🔗 Example

Rebuild all objects with remote dependencies by executing the following procedure:

```
CALL CHECK_CATALOG ('REBUILD', null, null, null)
```

Repair all objects in <schema> by executing the following procedure:

```
CALL CHECK_CATALOG ('REPAIR', '<schema>', null, null)
```

Repair the object <schema>.<view> by executing the following procedure:

```
CALL CHECK_CATALOG ('REPAIR', '<schema>', '<view>', null)
```

Related Information

[Open Communication From Target to Source System \[page 1021\]](#)

[Set Up Trust Relationship Between Target and Source Systems \[page 1018\]](#)

[Create Credential for Authenticated Access to Source System \[page 1022\]](#)

[Enable and Configure Cross-Database Access \[page 219\]](#)

[Database-Specific Configuration Parameters \[page 293\]](#)

9.1.3 Copying a System using System Replication

SAP HANA system replication can be used to create a copy of an SAP HANA database in a quick and simple way.

You can register another SAP HANA database in one of the two system replication scenarios:

- As a secondary for a standalone SAP HANA database
- As a tier 3 secondary in a tier 2 system replication landscape

System Replication Scenarios

Original Setup	Source Database	Target Database
Standalone SAP HANA Database	Primary	Secondary
Tier 2 System Replication	Tier 2 Secondary	Tier 3 Secondary

After the replication is active and in sync, a takeover to the newly added tier makes the standalone SAP HANA database runnable with identical data as the source database.

9.1.3.1 Copy a System using System Replication

You can use SAP HANA system replication to create a copy of an SAP HANA database.

Prerequisites

- You need an SAP HANA database called source database, which is to be copied.
- You need a separate (virtual) host or more (virtual) hosts for the database copy called target database.

Procedure

1. Install an SAP HANA database of the same or a higher revision as the target database on the separate (virtual) host or on the (virtual) hosts.
2. Prepare the source database for replication using `hdbnsutil -sr_enable [--name=<siteName>]`
3. Register the target database either as a secondary or tier 3 secondary depending on your original setup using `hdbnsutil -sr_register --remoteHost=<primary master host> --remoteInstance=<primary instance id> --replicationMode=[sync|syncmem|async] --name=<sitename> --operationMode=[delta_datashipping|logreplay]`

If the parameter is given, the operation mode is set. The default operation mode is logreplay.
4. Start the newly registered target database.
5. When the system replication is active and in sync, perform a takeover on the target database.
6. After the takeover is done, the target database is running as a copy of the source database.
7. To avoid confusion with the source databases, rename the <SID> and change the instance number using the tool `hdblcm`.

9.1.4 Renaming a System

An SAP HANA system can be renamed by changing the system identifiers, like host names, SID, and instance number. Changing system identifiers can be performed with the SAP HANA database lifecycle manager (HDBLCM).

System Identifiers

System identifiers are required parameters set during SAP HANA system installation. In some cases, it is necessary to change the originally configured system identifiers. All three system identifiers - host name, SID, and instance number - can be changed together or individually from the SAP HANA database lifecycle manager graphical user or command-line interface.

i Note

System replication must be disabled before renaming the SAP HANA system. If you need to change system identifiers for a system that is set up for system replication, you must first disable system replication, then change the system identifiers on each host and finally re-enable system replication.

The following options are available for SAP HANA database lifecycle manager in graphical user and command-line interfaces:

Task	Graphical User Interface	Command-Line Interface
Rename an SAP HANA System Host	<ul style="list-style-type: none"> ▶ Rename the SAP HANA System ▶ Define Host Properties ▶ Edit Host ▶ 	<pre>--action=rename_system -- hostmap=<old host>=<new host></pre>
Change the SID of an SAP HANA System	<ul style="list-style-type: none"> ▶ Rename the SAP HANA System ▶ Define System Properties ▶ Target System ID ▶ 	<pre>--action=rename_system -- target_sid=<new sid></pre>
Change the Instance Number of an SAP HANA System	<ul style="list-style-type: none"> ▶ Rename the SAP HANA System ▶ Define System Properties ▶ Target Instance Number ▶ 	<pre>--action=rename_system -- number=<new instance number></pre>

Mounted SID Preparation

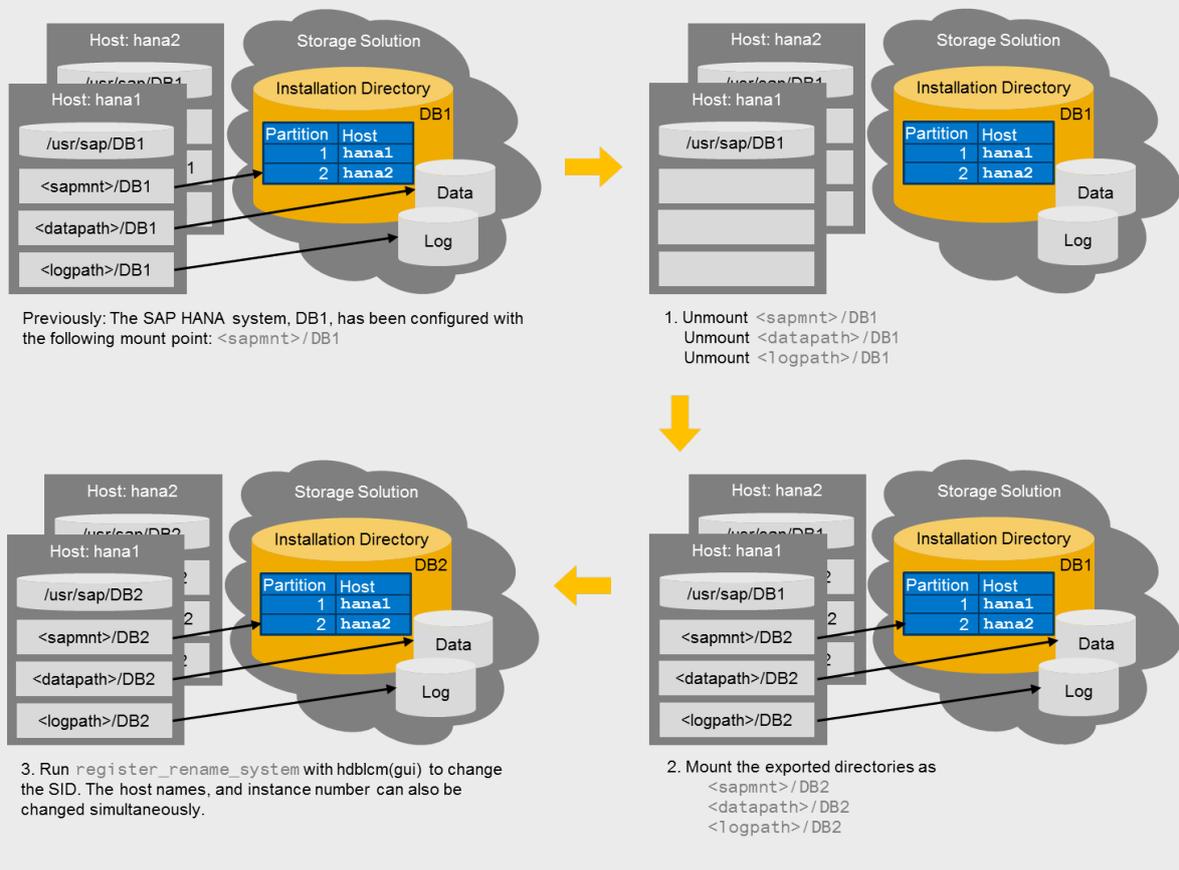
If the SID is included in the mount point, and you want to change the SID, you have to take extra preparation steps.

Normally, the installation path is exported and can be shared as `/<sapmnt>`. (The default for `<sapmnt>` is `/hana/shared`) so that several SAP HANA systems are located on the same physical device. However, if you exported a directory only for an individual SAP HANA system, the shared directory (the mount point) is `/<sapmnt>/<SID>`. In this case, you need to create a shared directory with the new target SID before changing the SID of the system.

❁ Example

In the following example, an SAP HANA system with a mounted SID is prepared for SID change:

Changing the SID for an SAP HANA System with a Mounted SID



9.1.4.1 Rename an SAP HANA System Host

You can rename an SAP HANA system host using SAP HANA database lifecycle manager (HDBLCM) resident program on the system which you want to configure.

Prerequisites

- You are logged in as root user.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- You are logged on as root user or as the system administrator user `<sid>adm`.
- The host you want to rename is either reachable via both the old and new host names or the SAP HANA system is stopped.
- The target SID must not exist. However, the target operating system administrator (`<SID>adm`) user may exist. Make sure that you have the password of the existing `<SID>adm` user, and that the user attributes

and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) resident program will not modify the properties of any existing user or group.

Context

i Note

If you rename an SAP HANA system, this usually invalidates the permanent SAP license. A temporary license is installed, and must be replaced within 28 days. For more information, see Related Information.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblc
```

By default, <sapmnt> is /hana/shared.

2. Start the rename task:

- To rename a host using the SAP HANA database lifecycle manager command-line interface:
 - Start the command-line tool interactively:

```
./hdblc
```

and enter the index of the `rename_system` action, or

- Start the tool with the `rename_system` action specified:

```
./hdblc --action=rename_system --hostmap=<old host>=<new host>
```

- To rename a host using the SAP HANA database lifecycle manager graphical user interface:
 1. Start the graphical user interface tool:

```
./hdblcgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

2. Choose *Rename the SAP HANA System*.
3. Select a host and choose *Edit Host...*
4. Enter the new host name in the *Target Host Name* field.

3. Define the required parameters.

For more information about parameters for the `rename_system` action, see the *SAP HANA Server Installation and Update Guide*.

i Note

When using the command line, the options can be set interactively during configuration only if they are marked as interactive in the help description. All other options have to be specified in the command

line. To call the help, in the SAP HANA resident HDBLCM directory of the SAP HANA system, execute the following command:

```
./hdbclm --action=rename_system --help
```

4. To continue with the task proceed as follows:
 - In the command line interface: Enter *y*.
 - In the graphical interface:
 1. *Next*.
 2. To execute the configuration task, choose To display the summary of the configuration data, choose *Rename* To display the summary of the configuration data, choose. The system displays the configuration progress.
 3. After the configuration task has finished, you can:
 - View the log. To do so, choose *View Log*.
 - Exit the graphical user interface. To do so, choose *Finish*.

Related Information

[Managing SAP HANA Licenses \[page 305\]](#)

9.1.4.2 Change the SID of an SAP HANA System

You can change the SID of an SAP HANA system using SAP HANA database lifecycle manager (HDBLCM) resident program on the system which you want to configure.

Prerequisites

- You are logged in as root user.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- The target SID must not exist. However, the target operating system administrator (<SID>adm) user may exist. Make sure that you have the password of the existing <SID>adm user, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) resident program will not modify the properties of any existing user or group.

Context

i Note

If you rename an SAP HANA system, this usually invalidates the permanent SAP license. A temporary license is installed, and must be replaced within 28 days.

An SAP HANA system has one SID for the system database and all tenants. Renaming a system changes the SID for the system database and all tenants.

Procedure

1. **In some cases**, the shared directory (mount point) includes the SID. If your mount point includes the SID, create a new shared directory before renaming the host.

Normally, the installation path (`<sapmnt>`), the data path (`<datapath>`), and the log path (`<logpath>`) are exported and can be shared. However, if you exported shared directories only for an individual SAP HANA system, the mount points are `<sapmnt>/<current SID>`, `<datapath>/<current SID>`, and `<logpath>/<current SID>`. In this case, you need to mount the exported directories as `<sapmnt>/<target SID>`, `<datapath>/<target SID>`, and `<logpath>/<target SID>` before changing the SID of the system.

- a. Stop the SAP HANA system.

To do this, in the SAP Host Agent perform the following operation:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function StopSystem
```

- b. Stop the sapstartsrv service by using the following SAP Host Agent operation:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function StopService
```

- c. Unmount `<sapmnt>/<current SID>`, `<datapath>/<current SID>`, `<logpath>/<current SID>`.
- d. Mount the exported directories as `<sapmnt>/<target SID>`, `<datapath>/<target SID>`, `<logpath>/<target SID>`.

2. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblcmlcm
```

By default, `<sapmnt>` is `/hana/shared`.

i Note

If the mount point includes the SID, change to the SAP HANA resident HDBLCM directory of the **target SID**.

3. Start the SID change task:
 - o To change the SID using the SAP HANA database lifecycle manager command-line interface:

- Start the command-line tool interactively:

```
./hdblcm
```

and enter the index of the `rename_system` action, or

- Start the tool with the `rename_system` action specified:

```
./hdblcm --action=rename_system --source_sid=<current SID> --target_sid=<new SID>
```

i Note

If the mount point includes the SID, and you have completed the preparation in Step 1, select the [Register and Rename SAP HANA System](#) action in either the SAP HANA database lifecycle manager graphical user interface or command-line interface.

- To change the SID using the SAP HANA database lifecycle manager graphical user interface:
 1. Start the graphical user interface tool:

```
./hdblcmgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

2. Choose [Rename SAP HANA System](#).
3. Enter the new SID in the [Target System ID](#) field.

4. Define the required parameters.

For more information about parameters for the `rename_system` and `register_rename_system` actions, see the *SAP HANA Server Installation and Update Guide*.

i Note

When using the command line, the options can be set interactively during configuration only if they are marked as interactive in the help description. All other options have to be specified in the command line. To call the help, in the SAP HANA resident HDBLCM directory of the SAP HANA system, execute the following command:

```
./hdblcm --action=rename_system --help
```

5. To continue with the task proceed as follows:
 - In the command line interface: Enter `y`.
 - In the graphical interface:
 1. [Next](#).
 2. To execute the configuration task, choose To display the summary of the configuration data, choose [Rename](#)To display the summary of the configuration data, choose. The system displays the configuration progress.
 3. After the configuration task has finished, you can:
 - View the log. To do so, choose [View Log](#).
 - Exit the graphical user interface. To do so, choose [Finish](#).

9.1.4.3 Change the Instance Number of an SAP HANA System

You can change the instance number of an SAP HANA system using SAP HANA database lifecycle manager (HDBLCM) resident program on the system which you want to configure.

Prerequisites

- You are logged in as root user.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.

Context

i Note

If you rename an SAP HANA system, this usually invalidates the permanent SAP license. A temporary license is installed, and must be replaced within 28 days. For more information, see Related Information.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblc
```

By default, <sapmnt> is /hana/shared.

2. Start the instance number change task:
 - To change an instance number using the SAP HANA database lifecycle manager command-line interface:
 - Start the command-line tool interactively:

```
./hdblc
```

and enter the index of the `rename_system` action, or

- Start the tool with the `rename_system` action specified:

```
./hdblc --action=rename_system --number=<new instance number>
```

- To rename a host using the SAP HANA database lifecycle manager graphical user interface:
 1. Start the graphical user interface tool:

```
./hdblcgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

2. Choose *Rename SAP HANA System*.

3. Define the required parameters.

For more information about parameters for the `rename_system` action, see the *SAP HANA Server Installation and Update Guide*.

i Note

When using the command line, the options can be set interactively during configuration only if they are marked as interactive in the help description. All other options have to be specified in the command line. To call the help, in the SAP HANA resident HDBLCM directory of the SAP HANA system, execute the following command:

```
./hdbclm --action=rename_system --help
```

4. To continue with the task proceed as follows:

- In the command line interface: Enter *y*.
- In the graphical interface:
 1. *Next*.
 2. To execute the configuration task, choose To display the summary of the configuration data, choose *Rename* To display the summary of the configuration data, choose. The system displays the configuration progress.
 3. After the configuration task has finished, you can:
 - View the log. To do so, choose *View Log*.
 - Exit the graphical user interface. To do so, choose *Finish*.

Related Information

[Managing SAP HANA Licenses \[page 305\]](#)

9.2 Network Administration

Set up your SAP HANA system and the corresponding data center and network configuration in line with your organization's environment and implementation considerations.

An SAP HANA data center deployment can range from a database running on a single host to a complex distributed system with multiple hosts located at a primary and one or more secondary sites, and supporting a distributed multi-terabyte database with full high availability and disaster recovery.

How you configure your network depends on a number of considerations, including:

- Support for traditional database clients, Web-based clients, and administrative connections
- The number of hosts used for the SAP HANA system, ranging from a single-host system to a complex distributed system with multiple hosts

- Support for high availability through the use of standby hosts, and support for disaster recovery through the use of multiple data centers
- Security and performance

SAP HANA has different types of network communication channels to support the different SAP HANA scenarios and setups:

- Channels used for external access to SAP HANA functionality by end-user clients, administration clients, application servers, and for data provisioning via SQL or HTTP
- Channels used for SAP HANA internal communication within the database or, in a distributed scenario, for communication between hosts

Before you start configuring the network for SAP HANA, it's important that you understand the different types of connections to, from, and within SAP HANA and which ports to configure for them. In addition, you should be familiar with the mechanisms used for assigning and resolving host names in SAP HANA.

Security

SAP HANA supports the isolation of internal communication from outside access. To separate external and internal communication, SAP HANA hosts use a separate network adapter with a separate IP address for each of the different networks. For IBM Power systems, this might be different.

In addition, SAP HANA can be configured to use TLS/SSL for secure communication. For more information, see the *SAP HANA Security Guide*.

Related Information

[Network Zones \[page 1041\]](#)

[Ports and Connections \[page 1043\]](#)

[Host Name Resolution \[page 1068\]](#)

[Configuring the Network for Multiple Hosts \[page 1438\]](#)

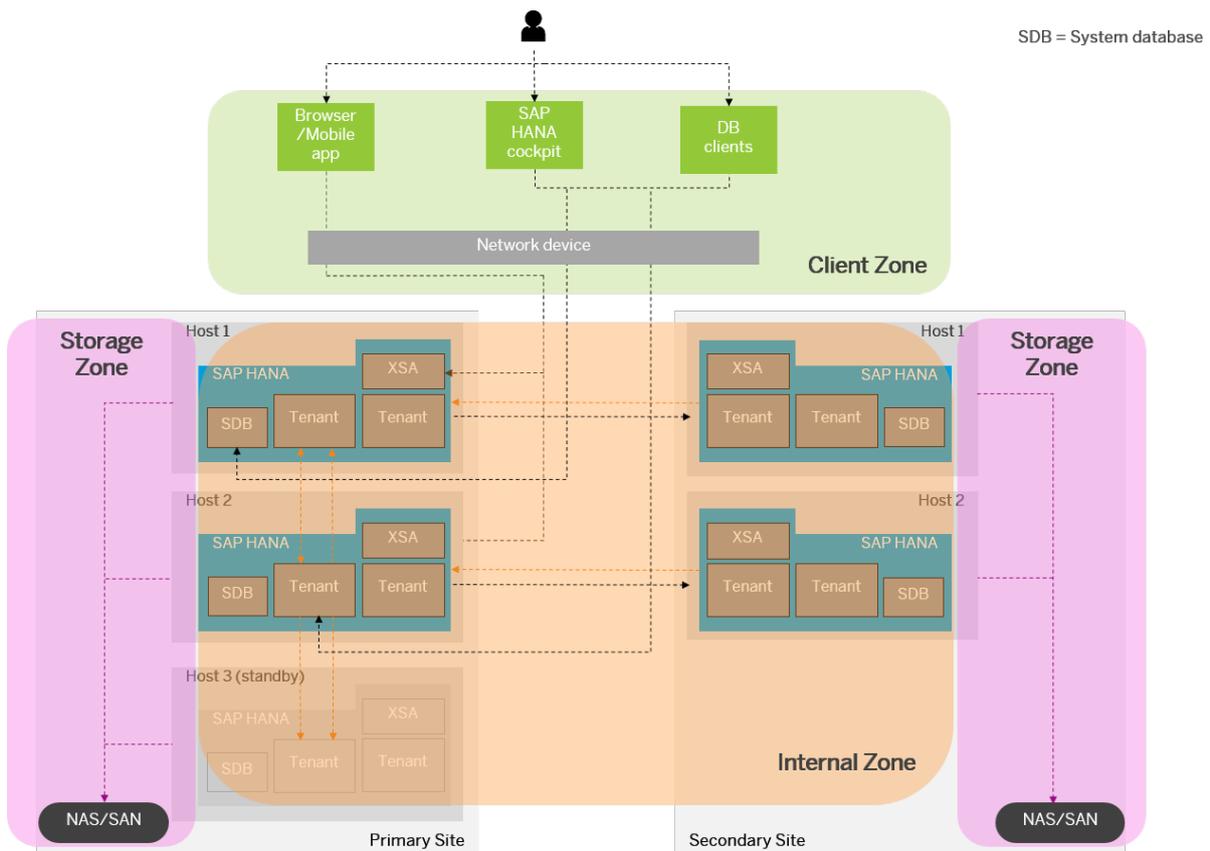
9.2.1 Network Zones

Separate network zones, each with its own configuration, allow you to control and limit network access to SAP HANA to only those channels required for your scenarios, while ensuring the required communication between all components in the SAP HANA network.

These network zones can be basically described as follows:

- Client zone
The network in this zone is used by SAP application servers, by clients such as the SAP HANA studio or Web applications running against the SAP HANA XS server, and by other data sources such as SAP Business Warehouse.

- Internal zone
This zone covers the interhost network between hosts in a distributed system as well as the SAP HANA system replication network.
- Storage zone
This zone refers to the network connections for backup storage and enterprise storage.
In most cases, the preferred storage solution involves separate, externally attached storage subsystem devices that are capable of providing dynamic mount-points for the different hosts, according to the overall landscape. A storage area network (SAN) can also be used for storage connectivity – for example, when running SAP HANA on IBM Power.



Related Information

[Connections from Database Clients and Web Clients to SAP HANA \[page 1043\]](#)

[Host Name Resolution for SQL Client Communication \[page 1074\]](#)

[Connections for Distributed SAP HANA Systems \[page 1053\]](#)

[Internal Host Name Resolution \[page 1072\]](#)

[Host Name Resolution for System Replication \[page 1119\]](#)

[SAP HANA - Storage Requirements](#)

[FAQ - SAP HANA Tailored Data Center Integration FAQ](#)

9.2.2 Ports and Connections

Before you start configuring the network for SAP HANA, you'll want to get an overview of the different types of connections to, from, and within SAP HANA and which ports to configure for them.

Related Information

[Connections from Database Clients and Web Clients to SAP HANA \[page 1043\]](#)

[Connections for Distributed SAP HANA Systems \[page 1053\]](#)

[Connections for SAP HANA Extended Application Services, Advanced Model \[page 1060\]](#)

[Connections for Components in the Extended SAP HANA Landscape \[page 1064\]](#)

9.2.2.1 Connections from Database Clients and Web Clients to SAP HANA

Several types of connections between SAP HANA and external clients are possible.

The connections between SAP HANA and external components and applications can be classified as follows:

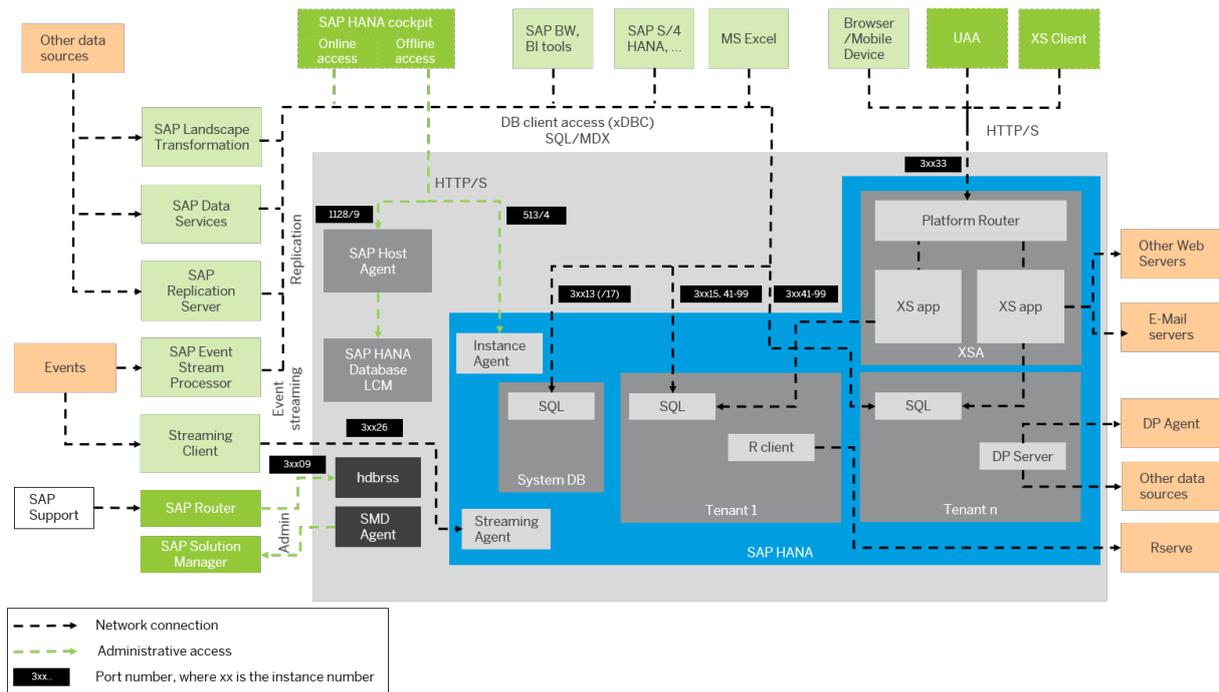
- Connections used for administrative purposes
- Connections used for data provisioning
- Connections from database clients that access the SQL/MDX interface of the SAP HANA database
- Connections from HTTP/S clients
- Outbound connections

You can see an example of what these connections look like in the diagram below. Network connections are depicted by dashed arrows. The direction of each arrow indicates which component is the initiator (start of arrow) and which component is the listener (end point of arrow). Administrative access to and from SAP HANA is depicted by the green dashed arrows. Port numbers are shown with a black background. The "xx" in the port numbers stands for the SAP HANA instance number.

The diagram shows all the network channels used by the SAP HANA software. For purposes of illustration, the diagram shows a single-host installation with two tenant databases. However, the connections shown apply equally to a distributed scenario.

i Note

In distributed scenarios, you must also ensure that every database client can connect to every host (not shown in the diagram). Moreover, additional network channels are required in distributed scenarios for communication between the different hosts of an SAP HANA system or between the different sites. For more information, see the section on connections for distributed systems.



Connections from Database Clients and Web Clients to SAP HANA

The following tables explain the diagram and the different categories described above in more detail.

Database Client Access via ODBC/JDBC (SQLDBC)

Client	Additional Information	TCP Port
Application servers that use SAP HANA as a database	You must enable SQL/MDX access for all database clients.	<ul style="list-style-type: none"> System database: 3xx13 First tenant database if automatically created during installation: 3xx15 Tenant databases: 3xx41–3xx98
Examples: SAP Business Warehouse and one or more components of SAP Business Suite	External and internal host names are mapped for the purposes of database client access. You can change the default mapping.	The port numbers of tenant databases are assigned automatically from the available port number range according to availability at the time the database is created. Administrators can also explicitly specify which port numbers to use when they create the tenant database. For more information, see the section on port assignment in tenant databases.
End-user clients that access the SAP HANA database directly		
Example: Microsoft Excel		

Client	Additional Information	TCP Port
SAP HANA cockpit and SAP HANA studio This connection is used for administrative purposes (for example, to access user data, configuration data or trace files).		i Note In a new system created with one initial tenant database or an upgraded single-container system, the first tenant database is accessible on port 3xx15.

For more information about JDBC/ODBC client connections to the SAP HANA database, see the *SAP HANA Client Interface Programming Reference*.

HTTP/S Client Access

Client	Additional Information	TCP Port
XS advanced application UI (browser, mobile, and so on)	Access from applications based on SAP HANA Extended Application Services, Advanced Model (SAP HANA XS advanced)	3xx33 (hostname routing)
XS User Account and Authentication service (browser, mobile, and so on)	Access from the application client to the xscontroller-managed Web Dispatcher (platform router) for the purposes of user authentication.	i Note The SAP HANA XS advanced application server supports two routing modes: port routing and hostname routing. As URLs in the hostname routing scenario are user friendly and there is only a single public port, this mode is recommended for production usage and is depicted in the figure above. If port routing is configured, all public ports of the platform router are exposed: 51000 – 51500, 3xx30, and 3xx32. For more information, see SAP Note 2245631.
<ul style="list-style-type: none"> • XSA command line client • Client library (Java) • One or more SAP HANA XS advanced model applications used, for example, for administrative and/or monitoring purposes 	Access to the xscontroller-managed Web Dispatcher for purposes of data access	

Client	Additional Information	TCP Port
XS classic application UI (browser, mobile, and so on)	<p>Access for applications based on SAP HANA Extended Application Services, classic model (SAP HANA XS classic).</p> <p>The SAP HANA platform itself has a number of Web applications that run on SAP HANA XS classic, for example, the SAP HANA Web-based Development Workbench and SAP HANA Application Lifecycle Management.</p>	80xx/43xx
	<p>i Note</p> <p>SAP HANA XS, classic is deprecated as of SAP HANA 2.0 SPS 02. For more information, see SAP Note 2465027.</p>	<p>i Note</p> <p>This port is not depicted in the graphic above.</p>
SAP HANA Direct Extractor Connection (DXC)	This connection is used for ETL-based data acquisition. For more information, see the <i>SAP HANA Direct Extractor Connection Implementation Guide</i> .	
UI toolkit for SAP HANA Info Access	This connection is used for the SAP HANA Info Access HTTP search service. SAP HANA Info Access provides UI building blocks for developing browser-based search apps on SAP HANA. For more information, see the <i>SAP HANA Search Developer Guide</i> .	
SAP HANA cockpit, SAP HANA studio	This is the connection to the SAP HANA database lifecycle manager via SAP Host Agent. For more information about the SAP HANA database lifecycle manager, see the section on SAP HANA Platform Lifecycle Management.	1128 1129 (SSL)

i Note

SAP HANA XS advanced server components and application instances also expose ports that are not designed for external communication. For more information about these, see *Connections for SAP HANA Extended Application Services, Advanced Model* and the *SAP HANA Security Guide*.

Administrative Tasks

Client	Protocol	Additional Information	TCP Port
SAP support	Internal SAP protocol	The connection is not active by default because it is required only in certain support cases. To find out how to open a support connection, see the section on getting support.	3xx09
SAP HANA cockpit, SAP HANA studio	SQLDBC (ODBC/JDBC)	The connection to the instance agent acts as an administrative channel for low-level access to the SAP HANA instance to allow features such as starting or stopping of the SAP HANA database.	5xx13 5xx14 (SSL)

Other administrative tasks, mainly database administration, use the SQL/MDX channel of the database.

Data Provisioning

Client	Protocol	Additional Information	TCP Port
Replication systems for external data sources	SQLDBC (ODBC/JDBC)	<p>The following replication technologies may be used:</p> <ul style="list-style-type: none"> • SAP Landscape Transformation (SLT) • SAP Data Services (DS) • SAP Replication Server (not included with all licensed editions of SAP HANA) 	<ul style="list-style-type: none"> • System database: 3xx13 • Tenant databases: 3xx41–3xx98 <p>The port numbers of tenant databases are assigned automatically from the available port number range according to availability at the time the database is created. Administrators can also explicitly specify which port numbers to use when they create the tenant database. For more information, see the section on port assignment in tenant databases.</p>
	HTTP/S	SAP HANA Direct Extractor Connection (DXC). This technology uses HTTP/S access via the SAP HANA XS classic server.	80xx/43xx

i Note

In a new system created with one initial tenant database or an upgraded single-container system, the first tenant database is accessible on port 3xx15.

i Note

This port is not depicted in the graphic above.

Client	Protocol	Additional Information	TCP Port
Streaming client	XML/RPC	This connection is used for SAP HANA Streaming Analytics (supported on Intel-based platforms only)	3xx26

⚠ Caution

SAP HANA Streaming Analytics is an SAP HANA option. Be aware that you need additional licenses for SAP HANA options. For more information, see [Important Disclaimer for Features in SAP HANA Platform \[page 1980\]](#).

Outbound Connections

Connection	Additional information
From the SAP Solution Manager diagnostics (SMD) agent to SAP Solution Manager	For information about how to install the SAP Solution Manager diagnostics agent, see SAP Note 1858920.
Calls from SAP HANA Extended Application Services to external servers	Examples: a Web server or an e-mail server (depends on what applications your company has deployed)
Smart data access from SAP HANA to external data sources for data federation purposes	For more information, see the section on smart data access.
From SAP HANA to the R environment	Only required for scenarios which use the R integration supported by SAP HANA. For more information, see the <i>SAP HANA R Integration Guide</i> .
From the data provisioning (DP) server of the SAP HANA database to the data provisioning agent and, depending on the type of adapter used, to the external data source(s)	<p>This connection is used for SAP HANA smart data integration in scenarios where SAP HANA is deployed on premise. For more information, see the <i>Installation and Configuration Guide for SAP HANA Smart Data Integration and SAP HANA Smart Data Quality</i>.</p> <p>SAP HANA with the DP server can run on IBM Power. However, the data provisioning agent needs to be hosted on an Intel machine. It is possible to connect between the two.</p>

⚠ Caution

SAP HANA smart data integration is an SAP HANA option. Be aware that you need additional licenses for SAP HANA options. For more information, see [Important Disclaimer for Features in SAP HANA Platform \[page 1980\]](#).

Related Information

[Connections from Database Clients and Web Clients to SAP HANA \[page 1043\]](#)

[Host Name Resolution for SQL Client Communication \[page 1074\]](#)

[Connections for SAP HANA Extended Application Services, Advanced Model \[page 1060\]](#)

[Connections for Components in the Extended SAP HANA Landscape \[page 1064\]](#)

[SAP Note 2245631](#)

[SAP Note 1858920](#)

[SAP Note 2465027](#)

9.2.2.1.1 Port Assignment in Tenant Databases

Every tenant database has its own ports and connections for internal and external communication.

Every tenant database has dedicated ports for SQL and internal communication. There is also a dedicated port for HTTP-based client communication via the SAP HANA XS classic server, which runs by default as an embedded service in the index server.

However, there are no standard port number assignments. Port numbers are assigned automatically from the available port number range according to availability at the time the database is created or a service is added. Administrators can also explicitly specify which port numbers to use when they create a tenant database or add a service.

The only exceptions to this are the tenant database that is automatically created when you **install** a single-tenant system and when you **convert** a single-container system to a tenant database system. This database retains the port numbers of the original single-container system: 3<instance>03 (internal communication), 3<instance>15 (SQL), and 3<instance>08 (HTTP via SAP HANA classic server). The ports of any subsequently added tenant database are automatically assigned according to availability at the time.

The default port number range for tenant databases is 3<instance>40–3<instance>99. This means that the maximum number of tenant databases that can be created per instance is 20. However, you can increase this by reserving the port numbers of further instances. In the cockpit, a dialog will prompt you to do this, or you can configure the property `[multidb] reserved_instance_numbers` in the `global.ini` file. The default value of this property is 0. If you change the value to 1, the port numbers of one further instance are available (for example, 30040–30199 if the first instance is 00). If you change it to 2, the port numbers of two further instances are available (for example, 30040–30299 if the first instance is 00). And so on.

i Note

The port number of the **system database** are fixed: 3<instance>01 (internal), 3<instance>13 (SQL), and 3<instance>14 (HTTP via XS classic server). If restricted SQL access is enabled, port 3<instance>17 is used for SQL requests to the system database.

Let's look at some simple examples.

❁ Example

Example 1:

You perform a default SAP HANA system installation. This results in the automatic creation of a single tenant database. This tenant database has the following port numbers: 3<instance>03 (internal communication), 3<instance>15 (SQL), 3<instance>08 (HTTP via XS classic server).

Then, you create two additional tenant databases. Each of these tenant databases is automatically assigned port numbers for the following connection types:

- Internal communication
- SQL
- HTTP (This is the port of the XS classic server embedded in the index server.)

The second database is assigned ports 3<instance>40–42, and the third 3<instance>43–45.

Example 2:

You install a new SAP HANA system without creating an initial tenant, by running the SAP HANA database lifecycle manager (HDBLCM) with the `create_initial_tenant` flag set to `off`. Then, you create three tenant databases. Each of these tenant databases is automatically assigned the following port numbers:

The first tenant database is assigned port numbers 3<instance>40–42, the second ports 3<instance>43–45, and the third 3<instance>46–48.

Example 3:

You install a new SAP HANA system without creating an initial tenant, by running the SAP HANA database lifecycle manager (HDBLCM) with the `create_initial_tenant` flag set to `off`. Then, you create a tenant database. The same port numbers as above are assigned: 3<instance>40 (internal communication), 3<instance>41 (SQL), and 3<instance>42 (HTTP via XS classic server). Next, you add a separate xsengine service to the first database. This service is automatically assigned the next three available port numbers: 3<instance>43–45. Finally, you create a second tenant database. This tenant database is automatically assigned the next three available port numbers: 3<instance>46–48.

Example 4:

You convert a single-container system to a tenant database system. This results in the automatic creation of one tenant database. This tenant database has the same port numbers as the original single-container system: 3<instance>03 (internal communication), 3<instance>15 (SQL), 3<instance>08 (HTTP via XS classic server). Then, you add a second indexserver to the tenant database. It is automatically assigned port numbers 3<instance>40–42. Finally, you create a second tenant database. It is automatically assigned ports the next three available port numbers: 3<instance>43–45.

i Note

All of the above examples refer to single-host systems and are based on automatic port number assignment.

You can determine the ports used by a particular tenant database by querying the `M_SERVICES` system view, either from the tenant database itself or from the system database.

- From the tenant database: `SELECT SERVICE_NAME, PORT, SQL_PORT, (PORT + 2) HTTP_PORT FROM SYS.M_SERVICES WHERE ((SERVICE_NAME='indexserver' and COORDINATOR_TYPE='MASTER') or (SERVICE_NAME='xsengine'))`
- From the system database: `SELECT DATABASE_NAME, SERVICE_NAME, PORT, SQL_PORT, (PORT + 2) HTTP_PORT FROM SYS_DATABASES.M_SERVICES WHERE DATABASE_NAME='<DBNAME>' and ((SERVICE_NAME='indexserver' and COORDINATOR_TYPE='MASTER') or (SERVICE_NAME='xsengine'))`

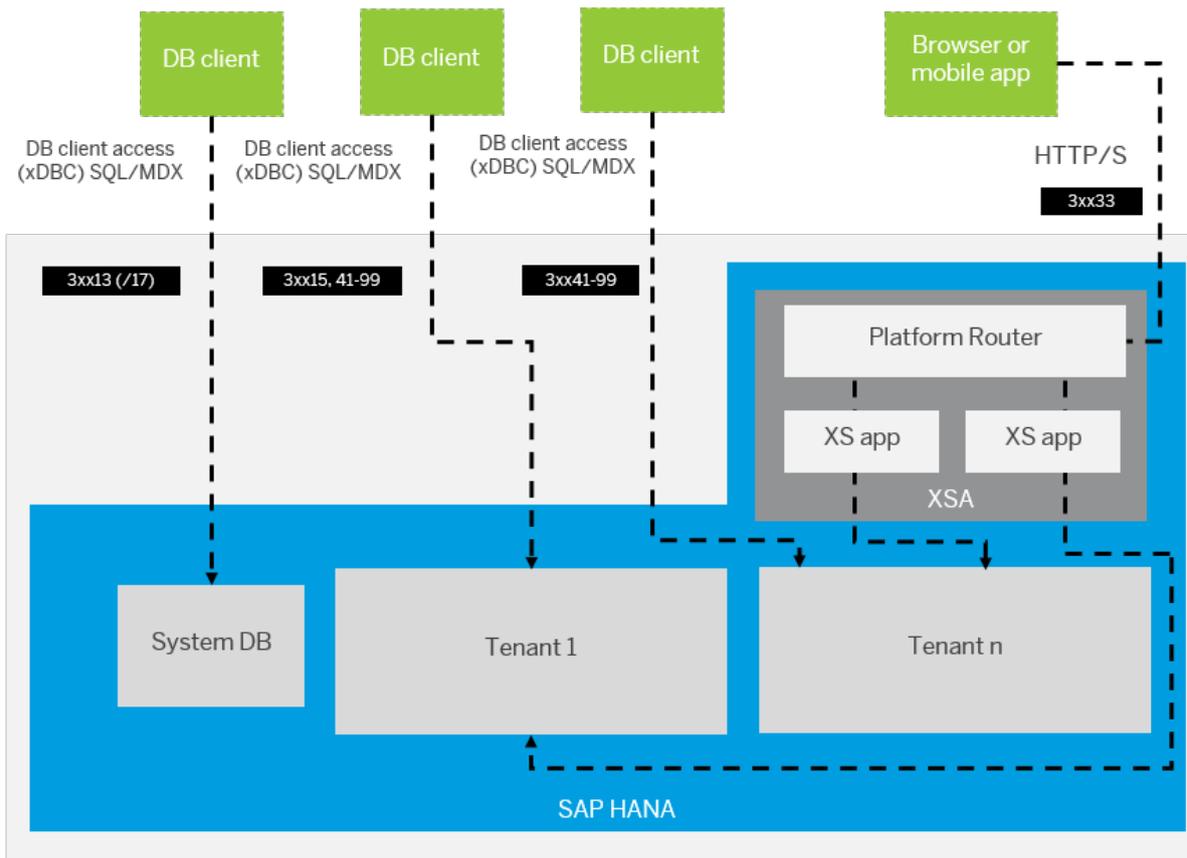
→ Remember

If your system was converted from single-container mode to a tenant database system, the HTTP port number of the first tenant database is **always** 3<instance>08 and not the port number returned using the above queries.

i Note

System privilege DATABASE ADMIN or CATALOG READ is required to read the M_SERVICES system view.

The following diagram shows an example of the connections and ports used in a system with two tenant databases, installed on a single host. The SAP HANA XS advanced runtime is used to run Web applications.



i Note

The SAP HANA XS advanced application server supports two routing modes: port routing and hostname routing. As URLs in the hostname routing scenario are user friendly and there is only a single public port, this mode is recommended for production usage and is depicted in the figure above. If port routing is

configured, all public ports of the platform router are exposed: 51000 – 51500, 3xx30, and 3xx32. For more information, see SAP Note 2245631.

Restricted SQL Access

If tenant databases need to be accessible from an external network, you can open an additional SQL port to prevent SQL access on port 3<instance>13. This prevents the exposure of the system database SQL administration port to the external network. You enable this feature by setting the property [multidb] `systemdb_separated_sql_port` to `true` in the `global.ini` file.

This opens port 17 for SQL requests to the system database and restricts access through port 3<instance>13 for database mapping. The connection through port 3<instance>13 is re-routed to 3<instance>17 if a connection to the system database is required. Make sure that port 17 is not exposed to the external network.

Related Information

[Creating and Configuring Tenant Databases \[page 189\]](#)

[SAP Note 2245631](#)

[Configure HTTP\(S\) Access to Tenant Databases via SAP HANA XS Classic \[page 1578\]](#)

9.2.2.2 Connections for Distributed SAP HANA Systems

SAP HANA systems can be distributed across multiple hosts for the purposes of scalability and availability.

An SAP HANA system is identified by a system ID (SID). Tenant databases are identified by the SID and their database name. Both are perceived as discrete units from the administration perspective. Some tasks are performed at system level (for example, installation, update, stop, and start) and others at database level (for example, database configuration, schema and table management). The system database is used for central system administration. The services of each database share the same metadata, and requests from client applications are transparently dispatched to the different services in the database. Servers that do not persist data, such as the compile server and the preprocessor server, run on the system database and serve all tenant databases.

A **distributed SAP HANA system** is a system that is installed on more than one host. An **SAP HANA instance** is a set of components of a distributed system that are installed on one host. Tenant databases can run individually on a single host or be distributed across several hosts.

Connections for Internal Communication Between Hosts and Sites

In addition to external network connections, SAP HANA uses separate, dedicated connections exclusively for internal communication. There are two types of internal communication in distributed systems:

- Communication between hosts in a multiple-host system (scale-out)
Internal network communication takes place between the hosts of a distributed system on one site. SAP HANA hosts contain a separate network interface card that is configured as part of a private network, using separate IP addresses and ports. For IBM Power systems, this might be different.

i Note

In single-host scenarios, the same communication channels are used for communication between the different processes on a single host and the internal IP addresses/ports are by default bound to the `localhost` interface.

There are a number of ways to isolate internal network ports from the client network. The preferred method depends on the data center configuration, on hardware vendor delivered options, and on the high availability implementation. Applying network separation for the internal communication prevents unauthorized access from outside networks. For additional security, it is possible to encrypt the internal communication using TLS/SSL. For more information, see the *SAP HANA Security Guide*.

Ports for Multiple-Host System

Client	TCP Port	Service	Note
Hosts of a distributed system on one site	3xx00	daemon	
	3xx01	nameserver	System database port only
	3xx02	preprocessor	System database port only
	3xx03	indexserver	Port used by the indexserver of the initial tenant database created in a new installation or an upgraded single-container system.
	3xx40 - 3xx97	indexserver	Tenant database ports Port numbers are assigned automatically from the available port number range according to availability at the time the database is created or a service is added. Administrators can also explicitly specify which port numbers to use when they create a tenant database or add a service.

Client	TCP Port	Service	Note
	3xx04	scriptserver	Port used by the script server of the first tenant database created in a new installation or an upgraded single-container system (optional)
	3xx40 - 3xx97	scriptserver	Tenant database ports (optional) Port numbers are assigned automatically from the available port number range according to availability at the time the database is created or a service is added. Administrators can also explicitly specify which port numbers to use when they create a tenant database or add a service.
	3xx40 - 3xx97	docstore	Tenant database ports (optional) Port numbers are assigned automatically from the available port number range according to availability at the time the database is created or a service is added. Administrators can also explicitly specify which port numbers to use when they create a tenant database or add a service.
	3xx10	compileserver	System database port only

Client	TCP Port	Service	Note
	3xx33	xscontroller	<p>System database port only, if hostname routing is configured</p> <p>The SAP HANA XS advanced application server supports two routing modes: port routing and hostname routing. As URLs in the hostname routing scenario are user friendly and there is only a single public port (3xx33), this mode is recommended for production usage and is depicted in the figure below. If port routing is configured, all public ports of the platform router are exposed: 51000 – 51500, 3xx30, and 3xx32.</p> <p>For more information, see SAP Note 2245631.</p>

i Note

The "xx" in the port numbers is the SAP HANA instance number.

- Communication between sites in a system replication configuration (high availability)
Internal network communication for system replication takes place between a primary site and a secondary site. In a multitier setup, this communication takes place between the tier-1 primary system and tier-2 secondary system as well as, asynchronously, between the tier-2 and tier-3 secondary systems. For more information about system replication and multitier setups, see the section on high availability. System replication is configured for the system as a whole. This means that the system database and all tenant databases are part of the system replication. System replication connections must be secured using TLS/SSL. In this case, landscape topology communication on the one hand, and data replication and log replication channels on the other, are secured in separate steps. For more information about configuring SSL for internal communication as well as securing communication between sites in system replication scenarios, see the *SAP HANA Security Guide*.

Ports for System Replication

Client	TCP Port	Service	Used For
Hosts on primary and secondary sites	4xx01	nameserver	Log and data shipping (system database)
	4xx02	nameserver	Unencrypted metadata communication (system database)

Client	TCP Port	Service	Used For
	4xx06	nameserver	TLS/SSL encrypted metadata communication (system database)
	4xx40 - 4xx97	indexserver	Log and data shipping (tenant databases)
	4xx40 - 4xx97	scriptserver	Log and data shipping (optional, tenant databases)
	4xx40 - 4xx97	docstore	Log and data shipping (optional, tenant databases)

i Note

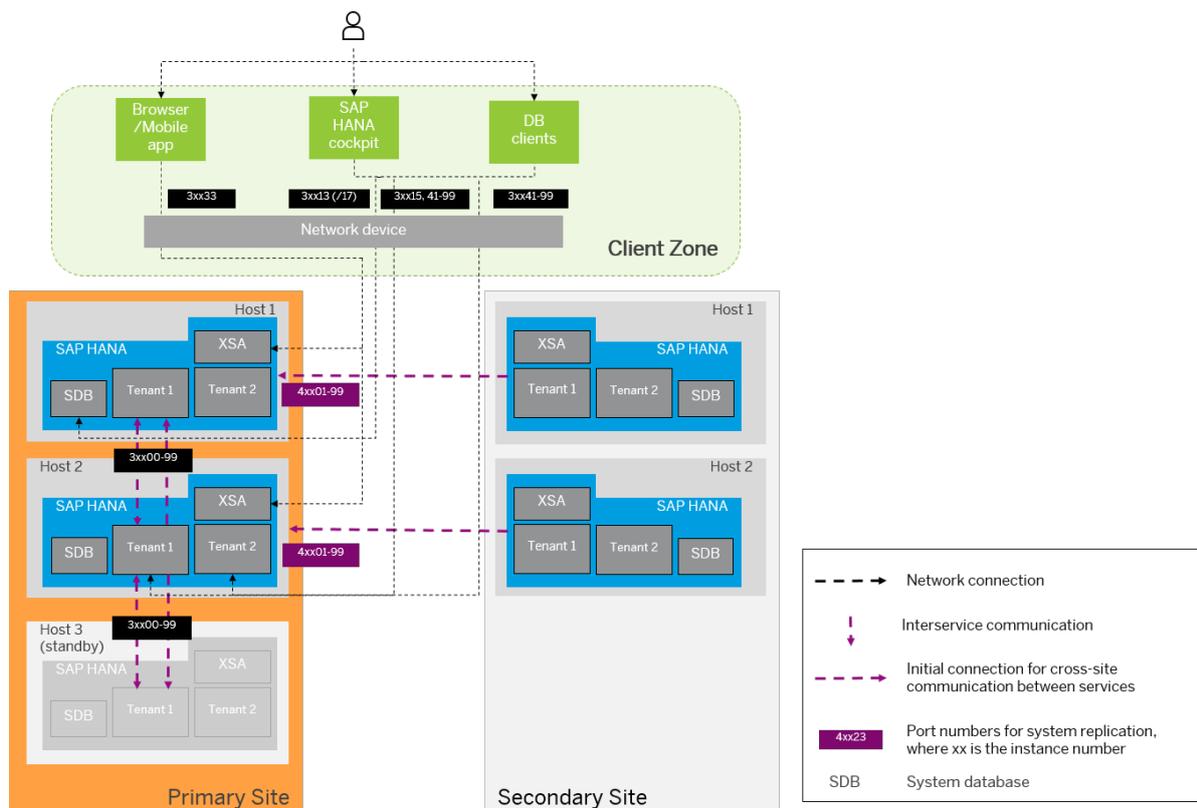
A port offset value of 10,000 is used to reserve ports for system replication communication. This shifts the ports from the 3<instance number>00 to the 4<instance number>00 port range for services.

i Note

SAP HANA internal communication has sometimes been unofficially referred to as TREXNet communication. However, the term TREXNet is not valid in the context of SAP HANA.

Example 1

The following diagram shows a multiple-host SAP HANA system with two active hosts and an extra standby host, fully system replicated to a secondary site to provide full disaster recovery support.



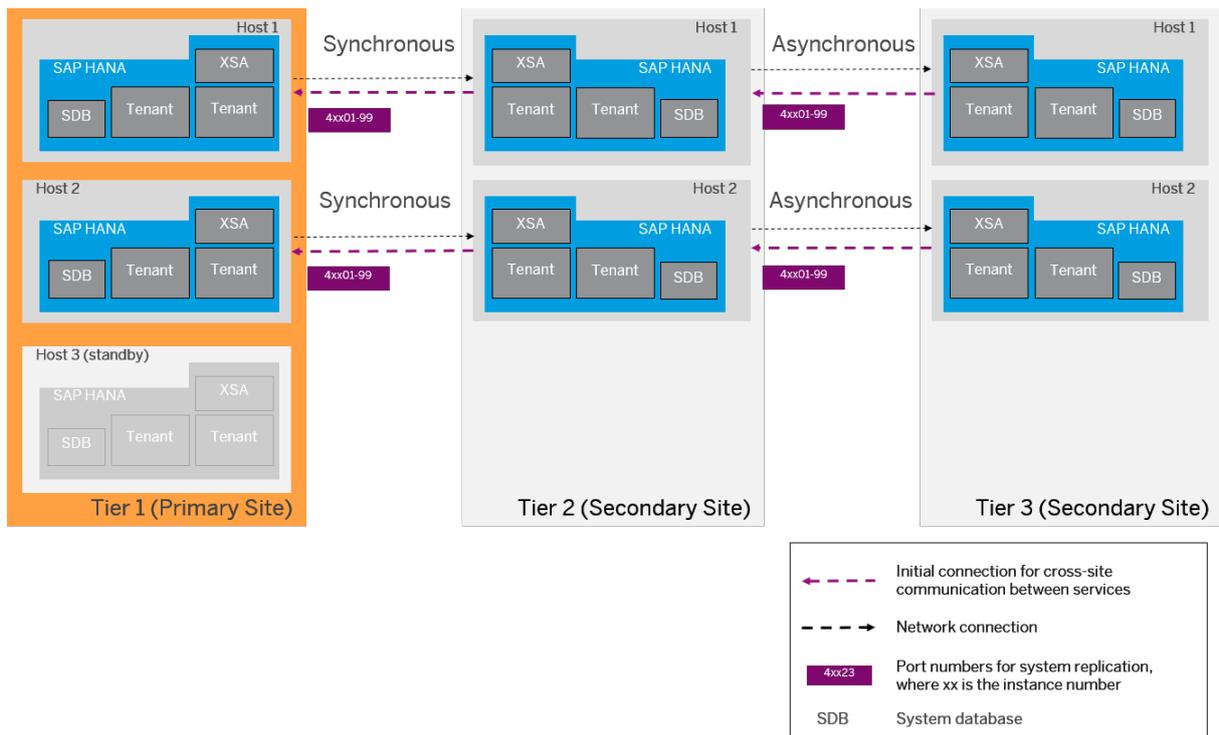
The vertical purple dashed lines show the communications between the services of a database; all instances communicate with all other instances of a multiple-host system on one site. The horizontal purple dashed lines show the initial connection for system replication communication between services on hosts on the primary site and the corresponding services on hosts of the secondary site (typically over a high-performance fiber network). The details of system replication configuration depend on the specific network setup of your company.

One of the most critical aspects of the network design of a highly available distributed system is the question of how the different clients manage to reconnect to the system when its topology changes due to the recovery operations following a failure or disaster. Two additional components can be used to handle client reconnection:

- A **network device** (router and/or switch), which can be used in conjunction with DNS or virtual IP redirection
- An **HTTP load balancer** (such as SAP Web Dispatcher) acts as a reverse proxy for HTTP connections and exposes a consistent external network address to the client network. The HTTP load balancer can also be used to provide load-balanced access to multiple distributed SAP HANA Extended Application Services (XS advanced) servers. For more information, see SAP Note 2300936.

Example 2

The following diagram shows an example of multi-tier system replication:



Related Information

[Internal Host Name Resolution \[page 1072\]](#)

[Connections from Database Clients and Web Clients to SAP HANA \[page 1043\]](#)

[Host Name Resolution for System Replication \[page 1119\]](#)

[SAP HANA System Replication with Tenant Databases \[page 1169\]](#)

[Host Auto-Failover Setup with XS Advanced Runtime \[page 1213\]](#)

[SAP Note 2245631](#)

[SAP Note 2300936](#)

9.2.2.3 Connections for SAP HANA Extended Application Services, Advanced Model

Additional ports and connections are required if you are using SAP HANA extended application services, advanced model (SAP HANA XS advanced).

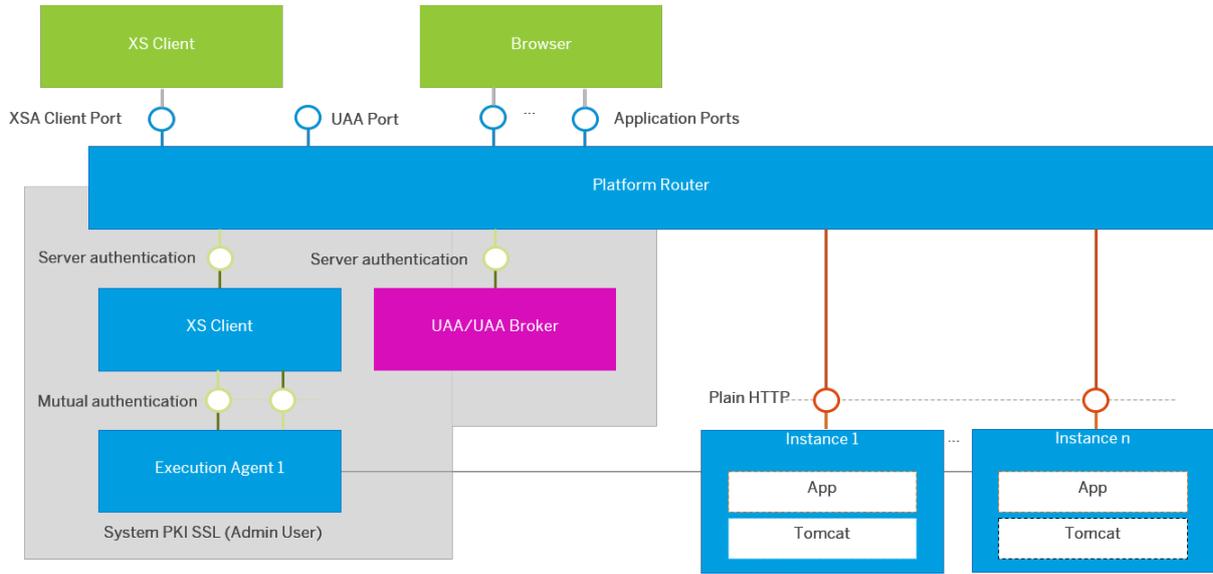
Client	Service	TCP Port	Use
Application UI (browser, mobile, and so on)	<code>xsuaaserver</code>	3xx32	Client HTTP(S) This port is used for the connection from the client to the <code>xscontroller-managed</code> Web Dispatcher (platform router) for purposes of user authentication.
Port 3xx32		3xx31	Internal HTTP(S) This port is used for the connection from the <code>xscontroller-managed</code> Web Dispatcher (platform router) to <code>xsuaaserver</code> for purposes of user authentication.
<ul style="list-style-type: none"> • Command line client • Client library (Java) • One or more SAP HANA XS advanced model applications used, for example, for administrative and/or monitoring purposes 	<code>xscontroller</code>	3xx30	Client HTTP(S) This port is used for the connection to the <code>xscontroller-managed</code> Web Dispatcher for purposes of data access.
Port 3xx30		Dynamic, in the range 51000-51500	Internal HTTP(S) This port range is used for the connection from the <code>xscontroller-managed</code> Web Dispatcher (platform router) to the <code>xscontroller</code> for purposes of data access.

Client	Service	TCP Port	Use
Application UI (browser, mobile, and so on)	Instances	Dynamic, in the range 51000-51500	Client HTTP(S) This port range is used for the connection from the client to the <code>xscontroller</code> -managed Web Dispatcher (Platform Router) for access to the application instance.
Application ports 51000-51500	Instances	Dynamic, in the range 50000-50999	Internal HTTP(S) This port range is used in single-host scenarios for the connection from the <code>xscontroller</code> -managed Web Dispatcher (platform router) to the application instances.
Application ports 51000-51500	Host-internal Web Dispatcher	Dynamic, in the range 50500-50999	Internal HTTP(S) This port range is used in multiple-host scenarios for the connection from the <code>xscontroller</code> -managed Web Dispatcher (Platform Router) to the host-internal platform router (host-specific Web Dispatcher).
Host-internal Web Dispatcher	Instances	Dynamic, in the range 50000-50499	Internal HTTP(S) This port range is used in multiple-host scenarios for the connection from the host-internal platform router (host-specific Web Dispatcher) to the application instance.
<code>xsexecagent</code>	<code>xscontroller</code>	3xx29	Internal HTTP(S)
<code>xscontroller</code>	<code>xsexecagent</code>	System	These ports are used for the connection between the <code>xs</code> execution agent and the <code>xscontroller</code> .

Client	Service	TCP Port	Use
Application UI (browser, mobile, and so on)	Application instances	3xx33	<p>Web Dispatcher HTTP(S)</p> <p>This port is used for the <code>xscontroller-managed</code> Web Dispatcher where routing is done by host names instead of ports. In this case, the <code>xscontroller</code> is available with URL <code>https://api.<example.com>:3xx33</code> and the <code>xsuaaserver</code> is available with URL <code>https://uaaserver.<example.com>:3xx33</code>.</p> <p>You specify the routing method – ports or host names – during installation. You can subsequently change the routing method in the <code>communication</code> section of the <code>xscontroller.ini</code> file.</p> <p>For more information, see SAP Note 2245631.</p>

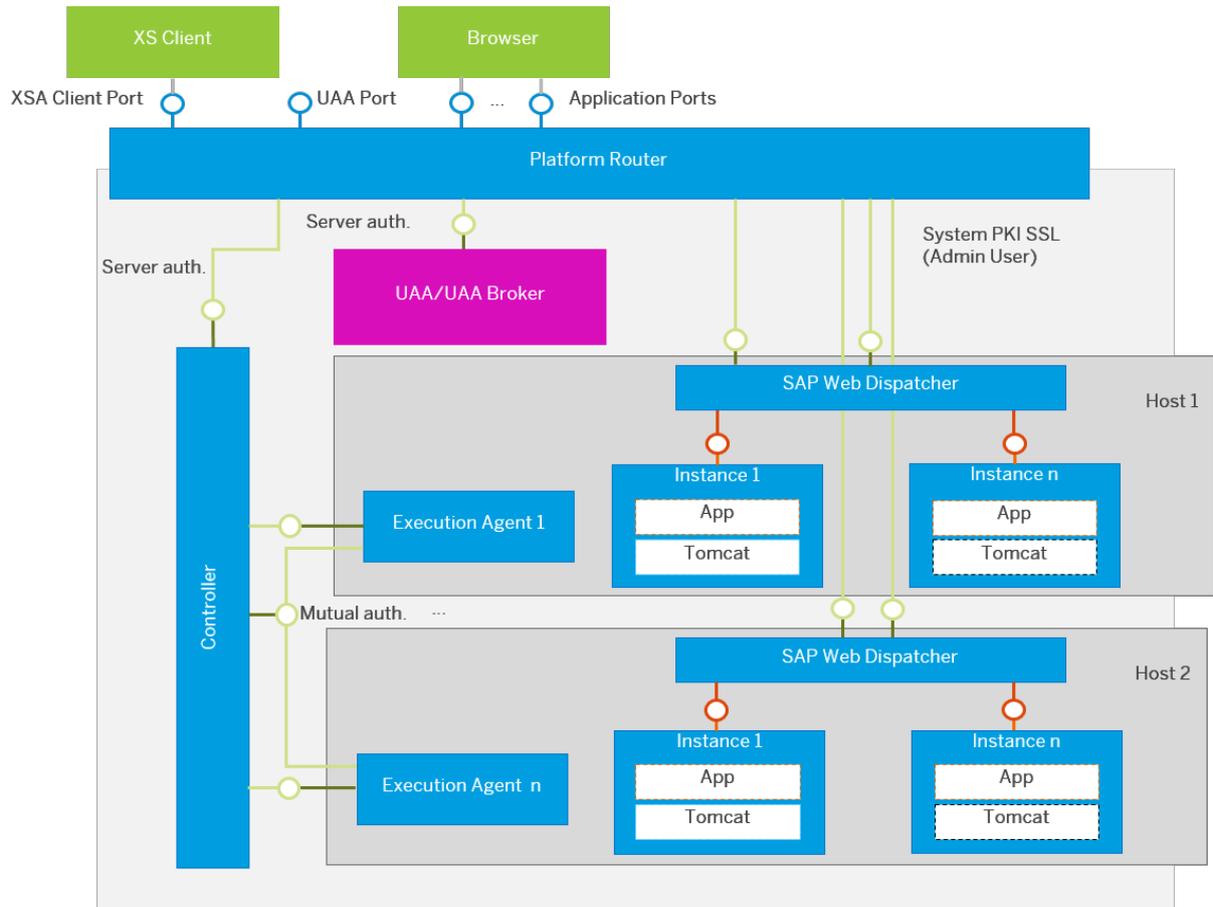
The following diagrams provide an overview of the ports and connections used by the SAP HANA XS advanced server infrastructure in single- and multiple-host scenarios. For more information about network and communication security, see the *SAP HANA Security Guide*.

XS Advanced Ports and Connections (Single-Host Scenario)



XS Advanced Ports and Connections (Single-Host Scenario)

XS Advanced Ports and Connections (Multiple-Host Scenario)



XS Advanced Ports and Connections (Multiple-Host Scenario)

Related Information

[SAP Note 2245631](#)

[SAP HANA System Architecture Overview \[page 17\]](#)

9.2.2.4 Connections for Components in the Extended SAP HANA Landscape

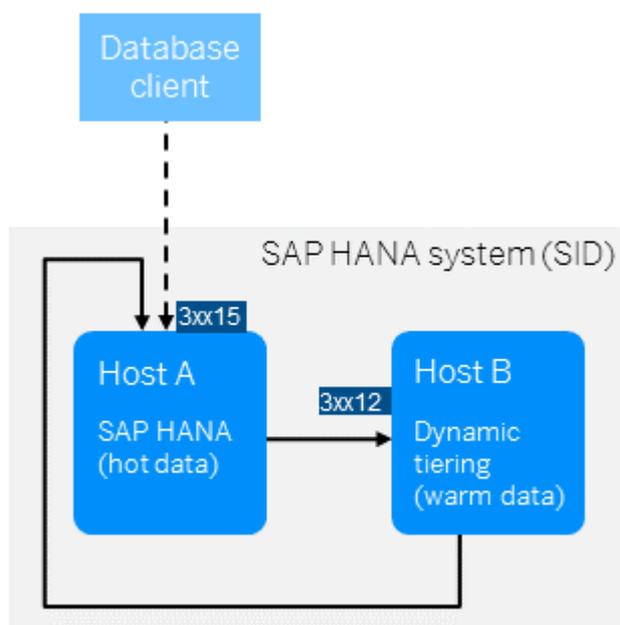
Dedicated ports are used to connect components in the extended SAP HANA landscape.

- [Connections for SAP HANA Dynamic Tiering \[page 1064\]](#)
- [Connections for SAP HANA Remote Data Sync \[page 1065\]](#)
- [Connections for SAP HANA Smart Data Integration \[page 1066\]](#)
- [Connections for SAP HANA Streaming Analytics \[page 1066\]](#)
- [Connections for SAP HANA Accelerator for SAP ASE \[page 1067\]](#)

9.2.2.4.1 Connections for SAP HANA Dynamic Tiering

No additional manual configuration of connections and ports is required in the SAP HANA software for SAP HANA dynamic tiering.

When an external client sends a request for warm data, it connects to the SAP HANA host which passes the request to the dynamic tiering host. The dynamic tiering host listens on internal port 3xx12. There is no direct connection between external components and the dynamic tiering host. The connection back from the dynamic tiering host to the SAP HANA host is through the SQL port of the SAP HANA host.



Related Information

[SAP HANA Dynamic Tiering](#)

9.2.2.4.2 Connections for SAP HANA Remote Data Sync

Connections between the components for SAP HANA remote data sync integration may differ depending on whether SAP HANA is deployed on premise, in the cloud, or behind a firewall.

Internal Connections for SAP HANA Remote Data Sync

The internal connections and ports for SAP HANA remote data sync are set up automatically. None of the ports are configurable.

i Note

SAP HANA remote data sync is supported on Linux x64 platforms only.

The SAP HANA server connects to one or more remote data sync servers on internal port 3xx27. Through this connection, SAP HANA gathers remote data sync statistics. The connection is triggered by the SAP HANA monitoring views.

The remote data sync hosts connect to the SQL port of the tenant database.

The remote data sync hosts retrieve the remote data sync license information. The connection is initiated upon start-up of the remote data sync host. When the remote data sync hosts are processing synchronization requests, many connections to SAP HANA are created and used to satisfy the requests. The maximum number of concurrent incoming requests to each remote data sync host is configured with the `max_network_connections` parameter in the `rdsyncserver.ini` file. The maximum number of concurrent database connections from each remote data sync host to the SAP HANA server is configured using the `db_worker_thread_count` parameter in the `rdsyncserver.ini` file.

During synchronization, if the script version being used has Java synchronization logic, then the configured Java scripts may also connect to SAP HANA and/or other external systems.

Synchronization clients run outside the SAP HANA system and connect to a remote data sync host on port 3xx28. One of several protocols are available on port 3xx28, including TCP, TLS, HTTP, and HTTPS. The protocol is configured with the `protocol` and `protocol_options` parameters in the `rdsyncserver.ini` file.

For more information, see the remote data sync documentation on SAP Help Portal.

Related Information

[SAP HANA Remote Data Sync](#)

9.2.2.4.3 Connections for SAP HANA Smart Data Integration

The connections between the components for SAP HANA smart data integration may differ depending on whether SAP HANA is deployed on premise, in the cloud, or behind a firewall.

For more information, see the *Installation and Configuration Guide* for SAP HANA Smart Data Integration and SAP HANA Smart Data Quality.

Related Information

[SAP HANA Smart Data Integration and SAP HANA Smart Data Quality](#)

9.2.2.4.4 Connections for SAP HANA Streaming Analytics

The internal connections and ports for SAP HANA streaming analytics are set up automatically. None of the ports are configurable.

i Note

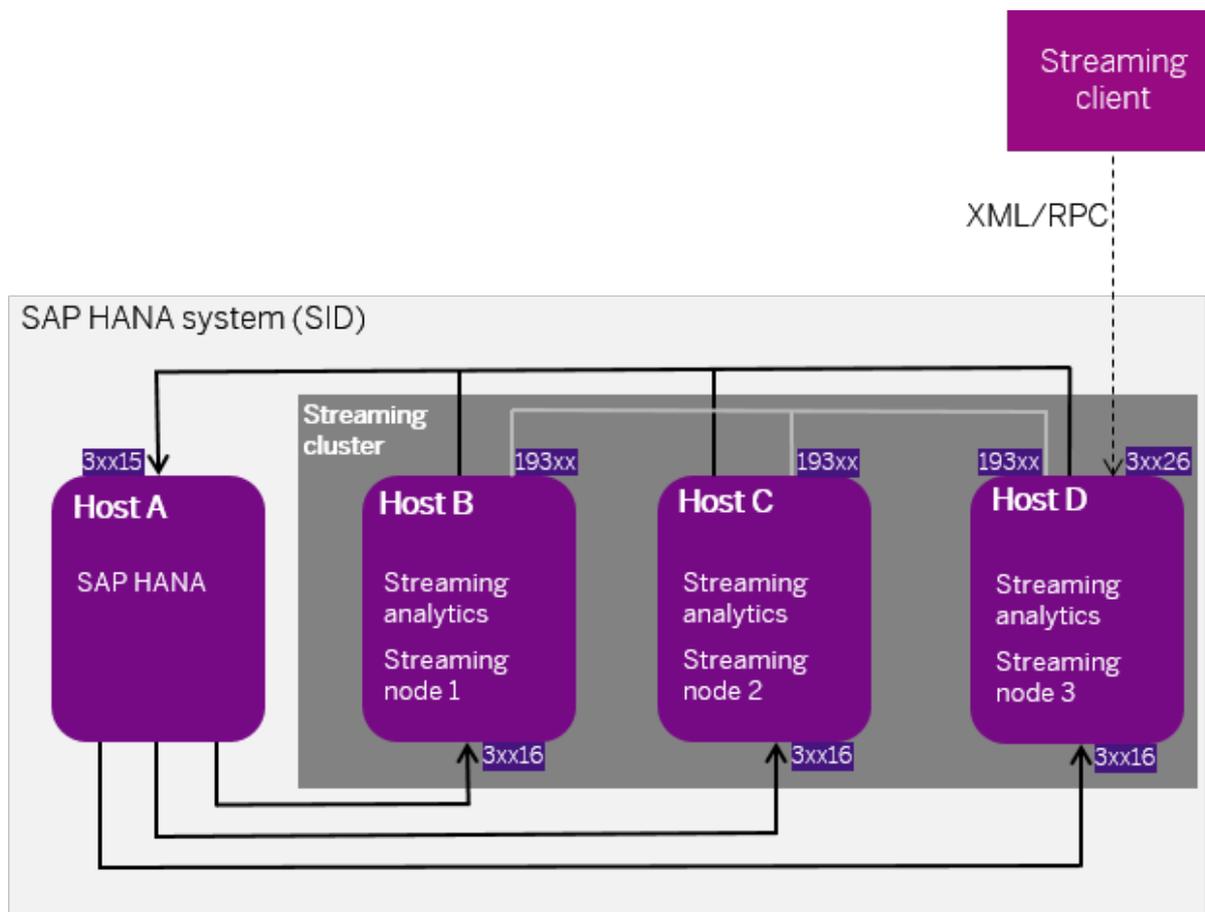
SAP HANA streaming analytics is supported on Intel-based platforms only.

The SAP HANA server connects to one or more streaming analytics servers on internal port 3xx16. Through this connection, SAP HANA gathers streaming analytics statistics. The connection is triggered by the SAP HANA cockpit monitoring views.

The streaming analytics hosts connect to the SAP HANA server on the SQL port of the tenant database. The streaming hosts retrieve the streaming license information and the streaming cluster configuration (which is stored on the SAP HANA database). If the streaming analytics project has an SAP HANA adapter or a generic database adapter that connects to SAP HANA, it also uses the SQL port connection.

Any streaming clients that run outside the SAP HANA system (such as custom-built external adapters) connect to a streaming node via the XML/RPC protocol on port 3xx26.

In a multi-node setup, the 193xx port is used for interserver communication between streaming hosts. This port is for internal use, but you may want to make a note of it for firewall settings.



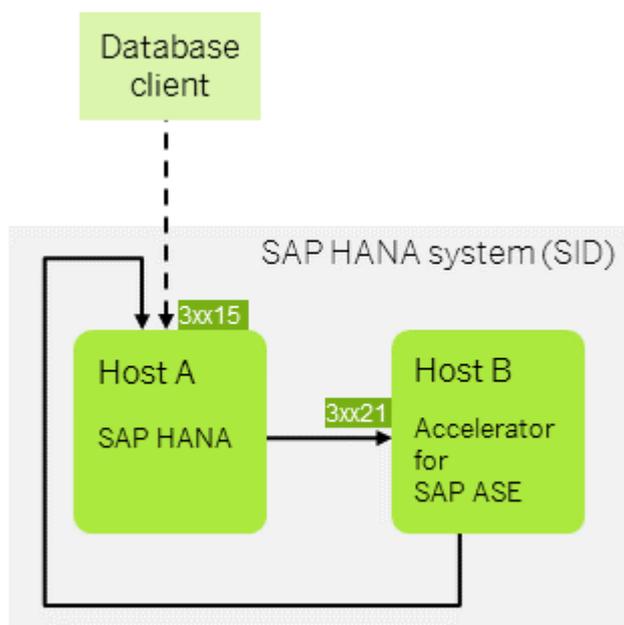
Related Information

[SAP HANA streaming analytics](#)

9.2.2.4.5 Connections for SAP HANA Accelerator for SAP ASE

The internal connections and ports for SAP HANA accelerator for SAP ASE are set up automatically.

When an external client sends a request for warm data, it connects to the SAP HANA host which passes the request to the accelerator for SAP ASE host. The accelerator for SAP ASE host listens on internal port 3xx21. The connection back from the accelerator for SAP ASE host to the SAP HANA host is through the SQL port of the SAP HANA host. Any SAP ASE clients that run outside the SAP HANA system can connect to an accelerator for SAP ASE node on port 3xx21 directly.



Related Information

[SAP HANA Accelerator for SAP ASE](#)

9.2.3 Host Name Resolution

Understand the mechanisms used for assigning and resolving host names in SAP HANA.

Related Information

[Default Host Names and Virtual Host Names \[page 1069\]](#)

[Internal Host Name Resolution \[page 1072\]](#)

[Host Name Resolution for System Replication \[page 1119\]](#)

[Host Name Resolution for SQL Client Communication \[page 1074\]](#)

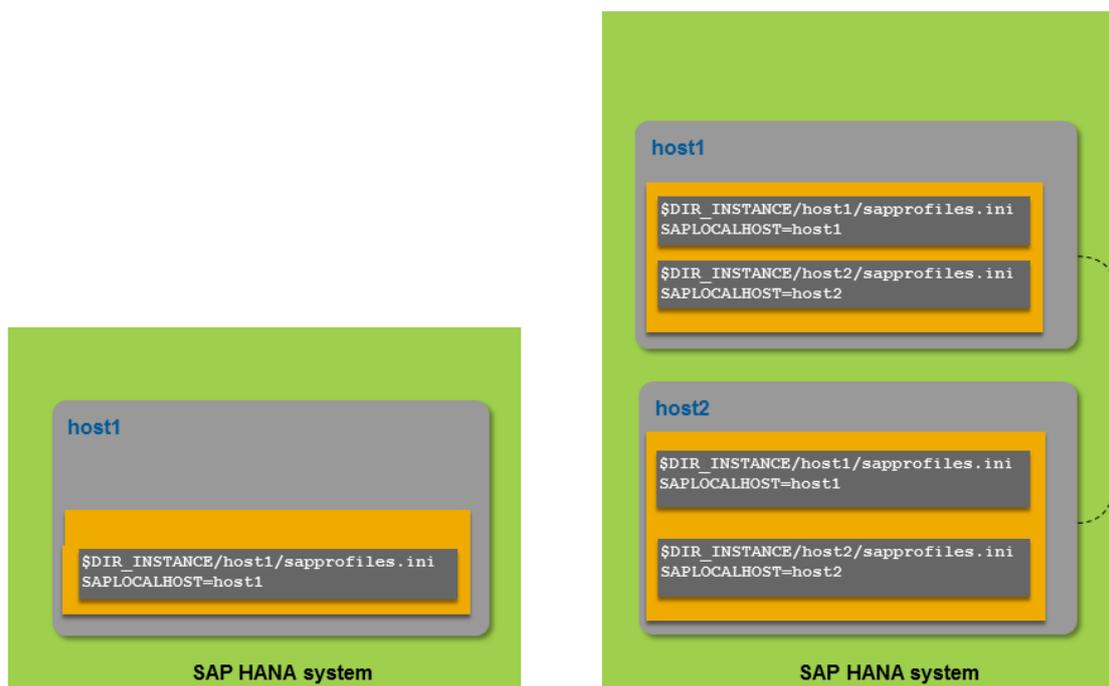
9.2.3.1 Default Host Names and Virtual Host Names

The assignment of multiple host names to the same host supports performance optimization as well as the security of your SAP HANA system. Moreover, some cluster managers and third-party backup tools as well as SAP Landscape Management work on the basis of virtual host names or IP aliases.

Default Host Names

The default host names, if nothing else is configured during the installation of SAP HANA, are the host names defined at operating system level. The installation extracts the host names known to the operating system (that is, the names of the SAP HANA instances) and stores them in the sapstart service profiles, that is, in the following files:

```
/usr/sap/sapservices  
/usr/sap/<SID>/HDB<instance_number>/<hostname>/sapprofile.ini
```



Example of Default Host Names for SAP HANA

These host names are then used for all internal communication between SAP HANA services (`nameserver`, `indexserver`, and so on) and the SAP start service (`sapstartsrv`). In addition, SAP HANA system views with a HOST column show these host names.

Virtual Host Names

Another approach is to specify alternative host names during installation. These are referred to as virtual host names. Virtual host names must also be unique across multiple SAP HANA systems if more than one data center or site is used.

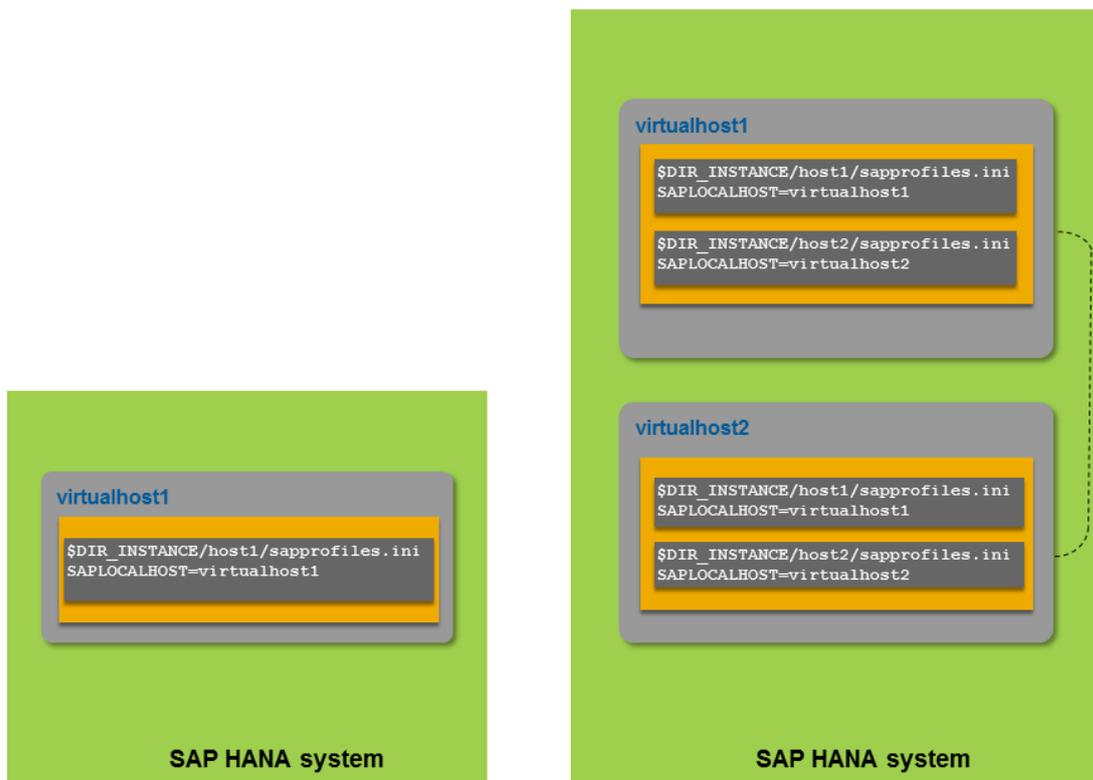
Host names specified in this manner must be resolvable during installation time, as well as when SAP HANA is in operation. This is achieved, for example, by adding an `<ip> <hostname>` line to the operating system file `/etc/hosts` that contains the hostname-to-IP address mappings for the TCP/IP subsystem. Here is an example of what this might look like at operating system level for one host:

```
127.0.0.1      localhost
10.68.91.226   virtualhost1.wdf.sap.corp virtualhost1
```

Virtual host names are assigned as part of the installation process with the platform LCM command-line tool `hdblcm` using the `hostname` parameter. For more information about using the command-line tool or the `hostname` parameter, see the *SAP HANA Server Installation and Update Guide*.

The `<virtualhostname>` is then stored as the internal host name in the `sapstart` service profiles and shows up in the `HOST` column of any system view.

It is also possible to assign virtual host names once the system is up and running by using the platform LCM action `system_rename` with the `hostmap` parameter. For more information about mapping hosts, see *Rename an SAP HANA System Host*.



Example of Virtual (Internal) Host Names for SAP HANA

Distributed Landscapes

In multiple-host systems used for scale-out, the host names of all hosts must be known to each SAP HANA host. The `/etc/hosts` file for each host must include the corresponding lines:

```

host1
127.0.0.1      localhost
10.68.91.226  virtualhost1.wdf.sap.corp virtualhost1
10.68.91.227  virtualhost2.wdf.sap.corp virtualhost2

```

```

host2
127.0.0.1      localhost
10.68.91.226  virtualhost1.wdf.sap.corp virtualhost1
10.68.91.227  virtualhost2.wdf.sap.corp virtualhost2

```

Related Information

[Use the Command-Line Interface to Perform Platform LCM Tasks \[page 924\]](#)

[Rename an SAP HANA System Host \[page 1034\]](#)

[Internal Host Name Resolution \[page 1072\]](#)

[Host Name Resolution for System Replication \[page 1119\]](#)

[Host Name Resolution for SQL Client Communication \[page 1074\]](#)

9.2.3.2 Internal Host Name Resolution

SAP HANA services use IP addresses to communicate with each other. Host names are mapped to these IP addresses through internal host name resolution, a technique by which the use of specific and/or fast networks can be enforced and communication restricted to a specific network.

Single Host Versus Multiple Hosts

For single-host systems, no additional configuration is required. The services listen on the loopback interface only (IP address 127.0.0.1). In the `global.ini` files, the `[communication] listeninterface` is set to `.local` as follows:

```
global.ini
[communication]
listeninterface=.local
```

In a distributed scenario with multiple hosts, the network needs to be configured so that interservice communication is operational throughout the entire landscape. In this setup, the host names (these could be virtual host names) of all hosts must be known to each other and thus to the SAP HANA system. This can be achieved by manually adding all hosts to each `/etc/hosts` file on the operating system of each host.

A distributed system can run with or without a separate network definition for interservice communication.

Distributed System Without a Separate Internal Network

If no separate network is defined for internal communication, SAP HANA services listen on all available network interfaces. In the `global.ini` file, the listening interface is set to `.global` as follows:

```
global.ini
[communication]
listeninterface=.global
```

Caution

If the `listeninterface` parameter is set to `.global`, we strongly recommend that you secure the SAP HANA servers with additional measures such as a firewall and/or TLS/SSL. Otherwise, the internal service ports of the system are exposed and can be used to attack SAP HANA.

Distributed System with a Separate Internal Network

A distributed system can be configured with a dedicated internal network in one of the following ways:

- At installation time, using the **HDBLCM command line option** as in the following example:

```
<installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64/hdblcm --  
internal_network=10.66.128.0/20
```

- Post installation, using the **resident HDBLCM** from the GUI, command-line, or Web user interface. The following example, in command-line mode, binds the processes to this address only and to all local host interfaces. This option requires an internal network address entry:

```
<sapmnt>/<SID>/hdblcm/hdblcm --action=configure_internal_network --  
listen_interface=internal --internal_address=10.66.8/21
```

For more information, see *Configuring SAP HANA Inter-Service Communication*.

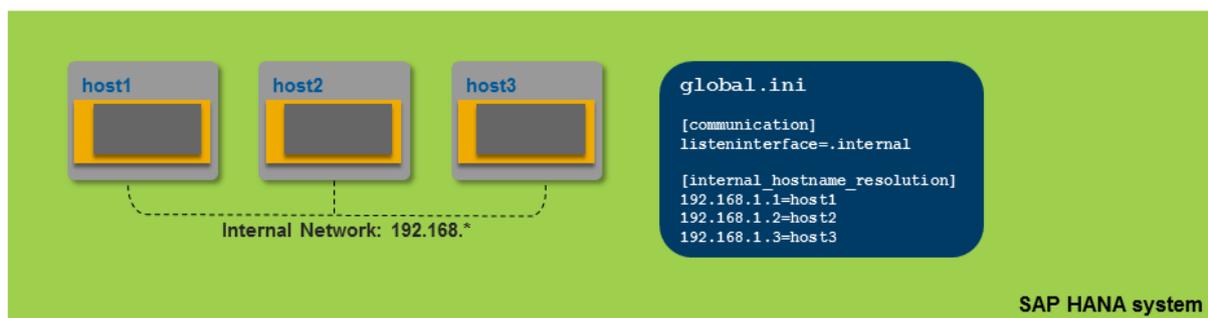
Post-installation configuration as described here is to be done by an SAP HANA system administrator with root credentials or on SAP HANA systems where SSH is configured. If root privileges or SSH are not available, you can still perform network configuration, but you need to use a host-by-host approach, also known as decentralized execution. In this case, see SAP Note 2048681.

SAP HANA automatically chooses on each host a network interface within the allowed network mask. If the network interface is defined as **.internal** in the `global.ini` file as described above, the SAP HANA services listen on this interface only:

```
global.ini  
[communication]  
listeninterface=.internal
```

Only the SAP start service (`sapstartsrv`) still listens on all interfaces, to accept start and stop commands, for example, from outside the SAP HANA system.

The following figure shows a simple example of how a separate internal network might be configured for an SAP HANA database with three hosts:



Simple Example of a Separate Internal Network for a Distributed SAP HANA System

For a more complex example, see *Host Name Resolution for System Replication*.

For more information about configuring the network for multiple hosts, see the section on scaling SAP HANA.

For information about the security of internal networks, see the *SAP HANA Security Guide*.

i Note

SAP HANA internal communication has sometimes been unofficially referred to as TREXNet communication. However, the term TREXNet is not valid in the context of SAP HANA.

Related Information

[Host Name Resolution for System Replication \[page 1119\]](#)

[Configuring the Network for Multiple Hosts \[page 1438\]](#)

[Configuring SAP HANA Inter-Service Communication \[page 1440\]](#)

[Configure SAP HANA Inter-Service Communication Using the Command-Line Interface \[page 1443\]](#)

[SAP Note 2048681](#) 

9.2.3.3 Host Name Resolution for SQL Client Communication

Client applications communicate with SAP HANA servers from different platforms and types of clients via a client library (such as SQLDBC, JDBC, ODBC, DBSL, ODBO or ADO.NET) for SQL or MDX access.

In distributed systems, the application has a **logical connection** to the SAP HANA system: that is, the client library may in fact use multiple connections to different servers or change to a different underlying connection. The client library supports load balancing and minimizes communication overhead by:

- Selecting connections based on load data
- Routing statements based on information about the location of data

i Note

Communication with SAP HANA hosts from a Web browser or a mobile application is requested using the HTTP protocol, which enables access to SAP HANA Extended Application Services, classic model (SAP HANA XS classic).

Public Host Name Resolution

An SQL client library always connects to the first available host specified in the connect string. From this host, the client library then receives a list of all the hosts. During operations, statements may be sent to any of these hosts.

By default, the IP address of the primary network interface is returned to the clients, as configured in the following parameter:

```
global.ini  
[public_hostname_resolution]
```

```
use_default_route=ip
```

This works as long as there is only one external network. If a hostname or IP address is unresolvable, the client library falls back on the host names in the connect string:

- In single-host systems, the user doesn't normally notice this. In rare cases, the connection attempt does not fail immediately but waits for a tcp timeout, making the first statement run very slowly.
- In distributed systems, performance is impaired because statements must first be sent to the initial host and then forwarded on the server side to the right host.

Connect String with Multiple Host Names

In a distributed SAP HANA system consisting of more than one host, a list of hosts (host:port) is specified in the SQL client library connect string.

The connect string for JDBC, for example, could look like this:

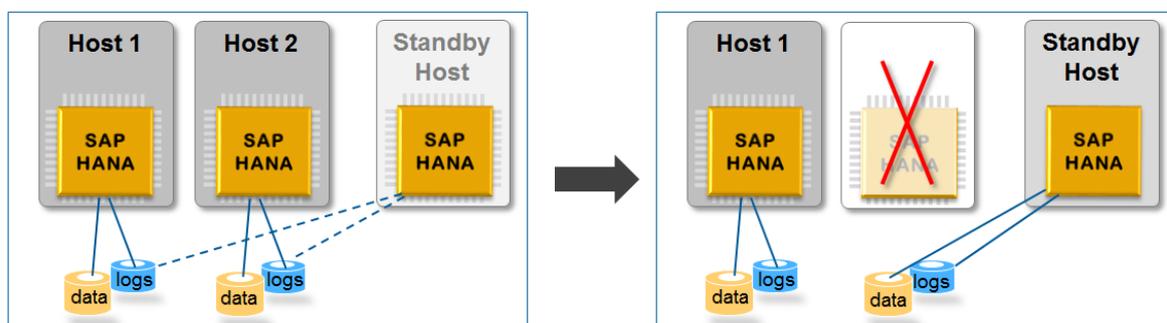
```
jdbc:sap://host1:30015;host2:30015;host3:30015/
```

All hosts that could become the active master, because they are one of the three configured master candidates, must be listed in the connect string to allow an initial connection to any of them in the event of a host auto-failover. A host auto-failover is an automatic switch from a crashed host to a standby host in the same system. One (or more) standby hosts are added to a SAP HANA system and configured to work in standby mode. As long as they are in standby mode, these hosts do not contain any data and do not accept requests or queries. When an active (worker) host fails, a standby host automatically takes its place.

Inclusion of the standby hosts in the connect string is mandatory if they are master candidates, otherwise optional.

The client connection code (ODBC/JDBC) uses a "round-robin" approach to reconnection, ensuring that the clients can always access the SAP HANA database, even after failover.

The following figure illustrates how host auto-failover works. An active host fails (in this example, Host 2), and the standby host takes over its role by starting its database instance using the persisted data and log files of the failed host.



Example of Auto Host-Failover

One way to look up the master candidates in your distributed SAP HANA database is to use the following SQL statement:

```
select HOST
from SYS.M_LANDSCAPE_HOST_CONFIGURATION
where NAMESERVER_CONFIG_ROLE like 'MASTER%'
order by NAMESERVER_CONFIG_ROLE
```

For more information, see the section on configuring clients for failover.

Connect String for SAP HANA System Replication

If system replication is used, we recommend that you do **not** specify physical host names in the SQL client connect string. Otherwise, you would have to reconfigure all of your applications after a takeover. Instead, use a **virtual host name** or **virtual IP address**, and manage it using an external cluster manager. This virtual host name or IP address must point to the active master host on the active primary site.

System replication takeover hooks can be implemented to provide notification about the takeover. For more information about takeover hooks and client connection recovery, see the section on system replication.

Related Information

[Mapping Host Names for Database Client Access \[page 1076\]](#)

[SQL Connection Information for New Clients \[page 1078\]](#)

[Configuring Clients for Failover \[page 1208\]](#)

[Client Connection Recovery after Takeover \[page 1132\]](#)

[SAP Note 1780950](#)

[SAP Note 1876398](#)

9.2.3.3.1 Mapping Host Names for Database Client Access

Clients communicate with the database through external host names or external IP addresses. A default mapping of external host names to internal host names enables statement routing and automatic reconnection in the event of a failover.

By default, the IP address of the primary network interface is used but there may be situations where you need to change this configuration, such as for certain firewall configurations, network address translation (NAT) types, or multiple external networks. For this purpose, a `[public_hostname_resolution]` section in the `global.ini` file is used with:

```
use_default_route = ip # values: no,ip,name,fqdn
optional pattern mapping: map_<internal-prefix>* = <public-prefix>*<public-
suffix>
optional exact mapping: map_<internal-name> = <public-name>
```

If optional mappings exist, they are always considered regardless of the `use_default_route` parameter value. Exact mappings have higher priority than pattern mappings.

Each host identifies the network interface and thus the default route for the connection:

Description	Parameter	Example
IP address of the interface	<code>use_default_route = ip</code>	10.4.2.71
Host name of the interface	<code>use_default_route = name</code>	lnd8520
Fully qualified name of the interface	<code>use_default_route = fqdn</code>	lnd8520.lnd.abc.corp
Disable feature and use internal host name	<code>use_default_route = no</code>	hananode01

❖ Example

For connections to tenant databases, certificate validation may not work due to how SAP HANA handles host name resolution. If this is the case, setting the value of the parameter `use_default_route` to `fqdn` on the SYSTEM layer ensures that SAP HANA uses the FQDN and that certificate validation for secured JDBC/ODBC connections is allowed.

In most cases, you do not need to configure anything. If you do need to configure something, see if you can use one of the default route mechanisms. You need to specify your own mapping only if the default route mechanisms do not fit your network requirements.

❖ Example

Here are some examples of how you might customize this parameter:

```
[public_hostname_resolution]
map_hananode* = myservername*
```

```
[public_hostname_resolution]
map_hananode* = hananode*.lnd.abc.corp
```

```
[public_hostname_resolution]
map_hananode01 = 10.4.2.71
map_hananode02 = 10.4.2.72
map_hananode03 = 10.4.2.73
map_hananode04 = 10.4.2.74
```

```
[public_hostname_resolution]
map_hananode0* = 10.4.2.7*
map_hananode1* = 10.4.2.8*
```

Changes to configuration and default routes are checked once a minute and become effective within a minute after the SQL system management statement `ALTER SYSTEM ALTER CONFIGURATION ... WITH RECONFIGURE .`

Related Information

[Host Name Resolution for SQL Client Communication \[page 1074\]](#)

9.2.3.3.2 SQL Connection Information for New Clients

It can be convenient for new SQL clients to be able to query the connectivity information of an existing client.

The connect string of the existing client was stored in the secure store and cannot be accessed. However, you can use the `global.ini/[communication]/sql_connect_hosts` parameter to record the connectivity information in the SAP HANA server so that it is available for the database connection from new clients. This information is a list of host names or IP addresses, which could be virtual host name or IP addresses, separated by commas.

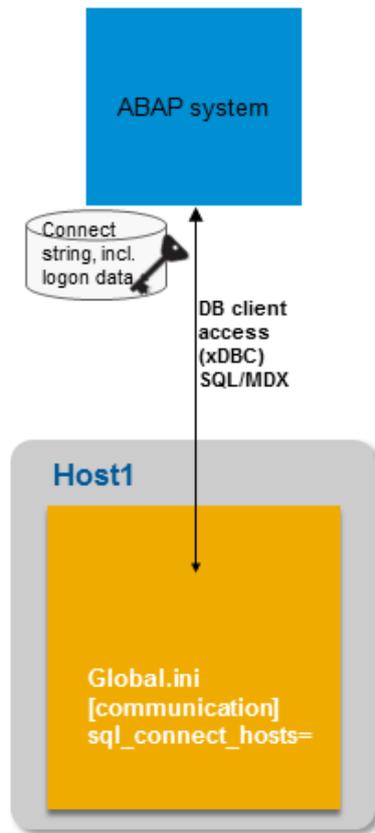
The SAP HANA server does not use this parameter. It is used by applications and components that connect to SAP HANA. If the parameter is not filled, the application needs to consume the host values as follows:

```
select HOST
from SYS.M_LANDSCAPE_HOST_CONFIGURATION
where NAMESERVER_CONFIG_ROLE like 'MASTER%'
order by NAMESERVER_CONFIG_ROLE
```

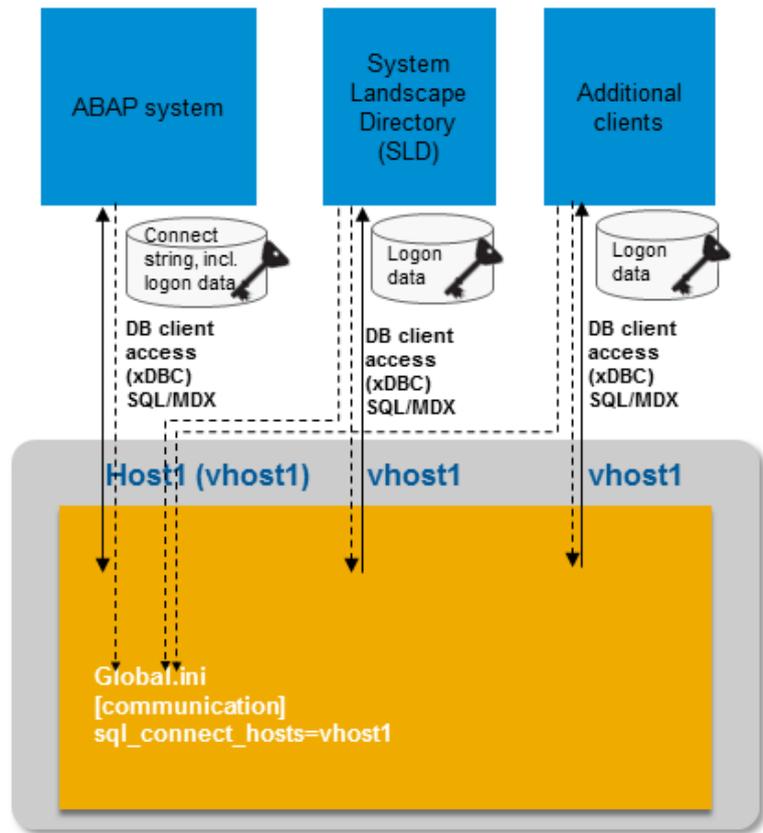
In the following example, an ABAP system is installed on SAP HANA and the connection information is stored on the client side in the connect string including the logon data. This is the standard case. The ABAP client then sets the `sql_connect_hosts` parameter on the SAP HANA server. A System Landscape Directory (SLD) is subsequently installed on the same SAP HANA system. The SLD agent is able to look up the parameter in SAP HANA to find out the connection information. If the parameter values are missing, SLD uses the above SQL statement. If more clients are added, they follow the same procedure.

The example shows a single host but the parameter can also be useful in scenarios with multiple hosts.

Initial SQL connection



With additional SQL connection



10 Availability and Scalability

SAP HANA provides comprehensive fault and disaster recovery support, as well as high availability for business continuity.

Related Information

[High Availability for SAP HANA \[page 1080\]](#)

[SAP HANA Database Backup and Recovery \[page 1229\]](#)

[Scaling SAP HANA \[page 1404\]](#)

10.1 High Availability for SAP HANA

High availability is the name given to a set of techniques, engineering practices, and design principles that support the goal of business continuity and also ensure that data and services are available to authorized users when needed.

SAP HANA is fully designed for high availability. It supports recovery measures ranging from faults and software errors, to disasters that decommission an entire data center. High availability is achieved by eliminating single points of failure (fault tolerance), and providing the ability to rapidly resume operations after a system outage with minimal business loss (fault resilience). Fault recovery is the process of recovering and resuming operations after an outage due to a fault. Disaster recovery is the process of recovering operations after an outage due to a prolonged data center or site failure. Preparing for disasters may require backing up data across longer distances, and may thus be more complex and costly.

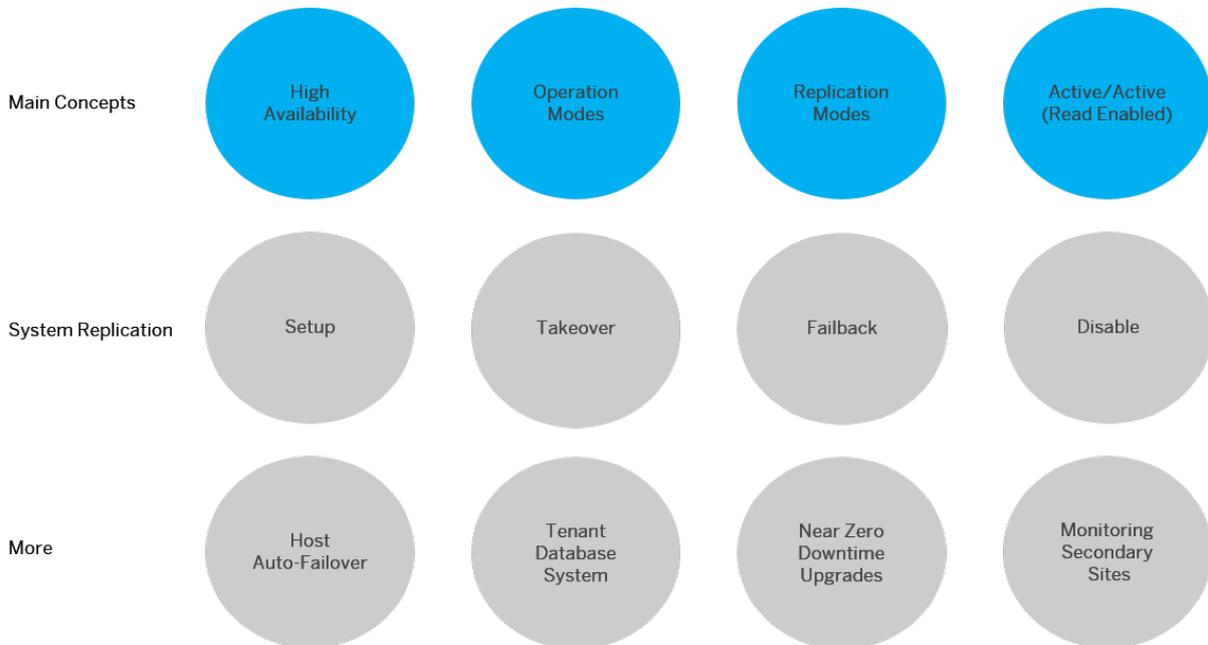
The key to achieving high availability is redundancy, including hardware redundancy, network redundancy and data center redundancy. SAP HANA provides several levels of defense against failure-related outages:

- **Hardware Redundancy:** SAP HANA appliance vendors offer multiple layers of redundant hardware, software and network components, such as redundant power supplies and fans, enterprise grade error-correcting memories, fully redundant network switches and routers, and uninterrupted power supply (UPS). Disk storage systems use batteries to guarantee writing even in the presence of power failure, and use striping and mirroring to provide redundancy for automatic recovery from disk failures. Generally speaking, all these redundancy solutions are transparent to SAP HANA's operation, but they form part of the defense against system outage due to single component failures.
- **Software:** SAP HANA is based on SUSE Linux Enterprise 11 for SAP and includes security pre-configurations (for example, minimal network services). Additionally, the SAP HANA system software also includes a watchdog function, which automatically restarts configured services (index server, name server, and so on), in case of detected stoppage (killed or crashed).
- **Persistence:** SAP HANA persists transaction logs, savepoints and snapshots to support system restart and recovery from host failures, with minimal delay and without loss of data.

- Standby and Failover: Separate, dedicated standby hosts are used for failover, in case of failure of the primary, active hosts. This improves the availability by significantly reducing the recovery time from an outage.

The diagram below helps you navigate through some of the most-searched topics on high availability.

Hover over each shape for a detailed description of the topic. Click the shape to open the topic.



- [SAP HANA High Availability Support \[page 1082\]](#)
- [Operation Modes for SAP HANA System Replication \[page 1094\]](#)
- [Replication Modes for SAP HANA System Replication \[page 1093\]](#)
- [Active/Active \(Read Enabled\) \[page 1157\]](#)
- [Setting Up SAP HANA System Replication \[page 1098\]](#)
- [Performing a Takeover \[page 1126\]](#)
- [Performing a Failback \[page 1136\]](#)
- [Disabling SAP HANA System Replication \[page 1141\]](#)
- [Setting Up Host Auto-Failover \[page 1208\]](#)
- [SAP HANA System Replication with Tenant Databases \[page 1169\]](#)
- [Use SAP HANA System Replication for Near Zero Downtime Upgrades \[page 1195\]](#)
- [Monitoring Secondary Sites \[page 1194\]](#)

Related Information

- [SAP HANA High Availability Support \[page 1082\]](#)
- [Operation Modes for SAP HANA System Replication \[page 1094\]](#)
- [Replication Modes for SAP HANA System Replication \[page 1093\]](#)
- [Active/Active \(Read Enabled\) \[page 1157\]](#)
- [Setting Up SAP HANA System Replication \[page 1098\]](#)

[Performing a Takeover \[page 1126\]](#)

[Performing a Failback \[page 1136\]](#)

[Disabling SAP HANA System Replication \[page 1141\]](#)

[Setting Up Host Auto-Failover \[page 1208\]](#)

[SAP HANA System Replication with Tenant Databases \[page 1169\]](#)

[Use SAP HANA System Replication for Near Zero Downtime Upgrades \[page 1195\]](#)

[Monitoring Secondary Sites \[page 1194\]](#)

10.1.1 SAP HANA High Availability Support

As an in-memory database, SAP HANA is not only concerned with maintaining the reliability of its data in the event of failures, but also with resuming operations with most of that data loaded back in memory as quickly as possible.

Downtime is the consequence of outages, which may be intentional (for example, for system upgrades) or caused by unplanned faults. A fault can be due to equipment malfunction, software or network failures, or due to a major disaster such as a fire, a regional power loss or a construction accident, which may decommission the entire data-center.

Fault Recovery is the process of recovering and resuming operations after an outage due to a fault. Disaster Recovery is the process of recovering operations after an outage due to a prolonged datacenter or site failure. Preparing for disasters may require backing up data across longer distances, and may thus be more complex and costly.

SAP HANA supports the following recovery measures from failures:

- Disaster recovery support:
 - Backups: Periodic saving of database copies in safe place.
 - Storage replication: Continuous replication (mirroring) between primary storage and backup storage over a network (may be synchronous).
 - System replication: Continuous update of secondary systems by primary system, including in-memory table loading.
- Fault recovery support:
 - Service auto-restart: Automatic restart of stopped services on host (watchdog).
 - Host auto-failover: Automatic failover from crashed host to standby host in the same system.
 - System replication: Continuous update of secondary systems by primary system, including in-memory table loading and read-only access on the secondary.

System replication is flexible enough that it can also be used for both fault and disaster recovery to achieve high availability. The data pre-load option can be used for fault recovery to enable a quicker takeover than with Host Auto-Failover. You can build a solution with single node systems and do not need a scale out system and the additional storage and associated costs.

SAP HANA supports system replication for tenant databases on the system level, this means the tenant database system as a whole including all tenant databases. An SAP HANA system installed in multiple-container mode always has exactly one system database and any number of tenant databases (including zero). For more information on tenant databases see *Creating and Configuring Tenant Databases*.

Using Secondary Servers for Non-Production systems

With SAP HANA system replication, you can use the servers on the secondary system for non-production SAP HANA systems under the following conditions:

- Table pre-load is turned off in the secondary system.
- The secondary system uses its own disk infrastructure. In the case of single node systems this means, the local disk infrastructure needs to be doubled.
- The non-production systems are stopped with the takeover to the production secondary.

Related Information

[Creating and Configuring Tenant Databases \[page 189\]](#)

[SAP Note 1999880](#)

[SAP Note 2183363](#)

[SAP Note 2300936](#)

SCN Documents

[SAP HANA Academy System Replication Videos](#)

[White paper "Introduction to High Availability for SAP HANA"](#)

[How to Perform System Replication for SAP HANA](#)

10.1.1.1 Backups

Backups are one of the key disaster recovery features offered by SAP HANA.

SAP HANA uses in-memory technology, but of course it fully persists any transaction that changes the data, such as row insertions, deletions and updates, so it can resume from a power-outage without loss of data. SAP HANA persists two types of data to storage: transaction redo logs, and data changes in the form of savepoints.

A transaction redo log is used to record a change. To make a transaction durable, it is not required to persist the complete data when the transaction is committed; instead it is sufficient to persist the redo log. Upon an outage, the most recent consistent state of the database can be restored by replaying the changes recorded in the log, redoing completed transactions and rolling back incomplete ones.

A savepoint is a periodic point in time, when all the changed data is written to storage, in the form of pages. One goal of performing savepoints is to speed up restart: when starting up the system, logs need not be processed from the beginning, but only from the last savepoint position. Savepoints are coordinated across all processes (called SAP HANA services) and instances of the database to ensure transaction consistency. By default, savepoints are performed every five minutes, but this can be configured.

Savepoints normally overwrite older savepoints, but it is possible to freeze a savepoint for future use; this is called a snapshot. Snapshots can be replicated in the form of full data backups, which can be used to restore a database to a specific point in time. This can be useful in the event of data corruption, for instance. In addition to data backups, smaller periodic log backups ensure the ability to recover from fatal storage faults with minimal loss of data.

Savepoints, can be saved to local storage, and the additional backups, can be additionally saved to backup storage. Local recovery from the crash uses the latest savepoint, and then replays the last logs, to recover the database without any data loss. If the local storage was corrupted by the crash, it is still possible to recover the database from the data and log backups, possibly with loss of some data. Regularly shipping backups to a remote location over a network or via couriers can be a simple and relatively inexpensive way to prepare for a disaster. Depending on the frequency and shipping method, this approach may have a recovery time ranging from hours to days.

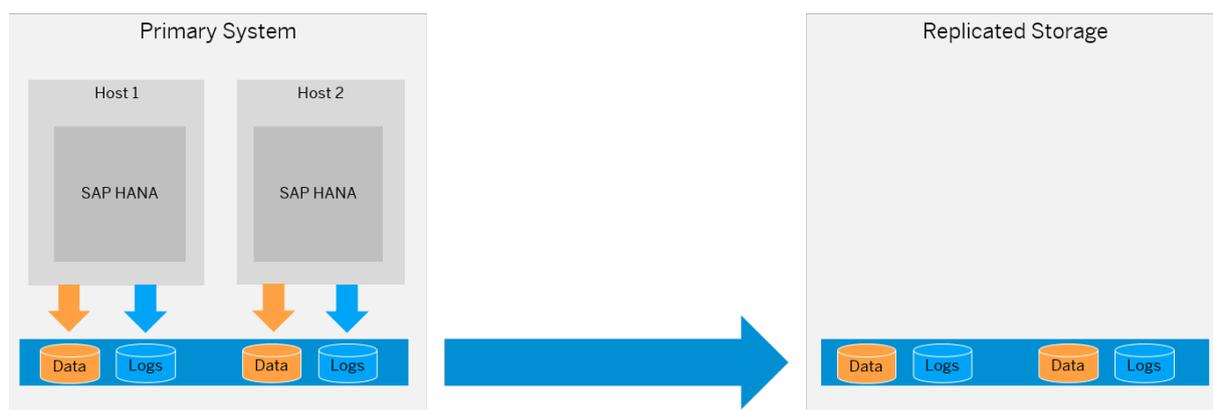
Related Information

[SAP HANA Database Backup and Recovery \[page 1229\]](#)

10.1.1.2 Storage Replication

SAP HANA offers disaster recovery support for storage replication solutions provided by hardware partners.

One drawback of backups is the potential loss of data between the time of the last backup and the time of the failure. A preferred solution therefore, is to provide continuous replication of all persisted data. Several SAP HANA hardware partners offer a storage-level replication solution, which delivers a backup of the volumes or file-system to a remote, networked storage system. In some of these vendor-specific solutions, which are certified by SAP, the SAP HANA transaction only completes when the locally persisted transaction log has been replicated remotely. This is called synchronous storage replication. Synchronous storage replication can be used only where the distance between the primary and backup site is relatively short (typically 100 kilometers or less), allowing for sub-millisecond round-trip latencies.



Due to its continuous nature, storage replication (sometimes also called remote storage mirroring) can be a more attractive option than backups, as it reduces the amount of time between the last backup and a failure. Another advantage of storage replication is that it also enables a much shorter recovery time. This solution requires a reliable, high bandwidth and low latency connection between the primary site and the secondary site.

See *SAP Note 1755396 Released DT solutions for SAP HANA with disk replication*

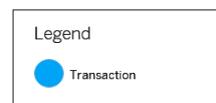
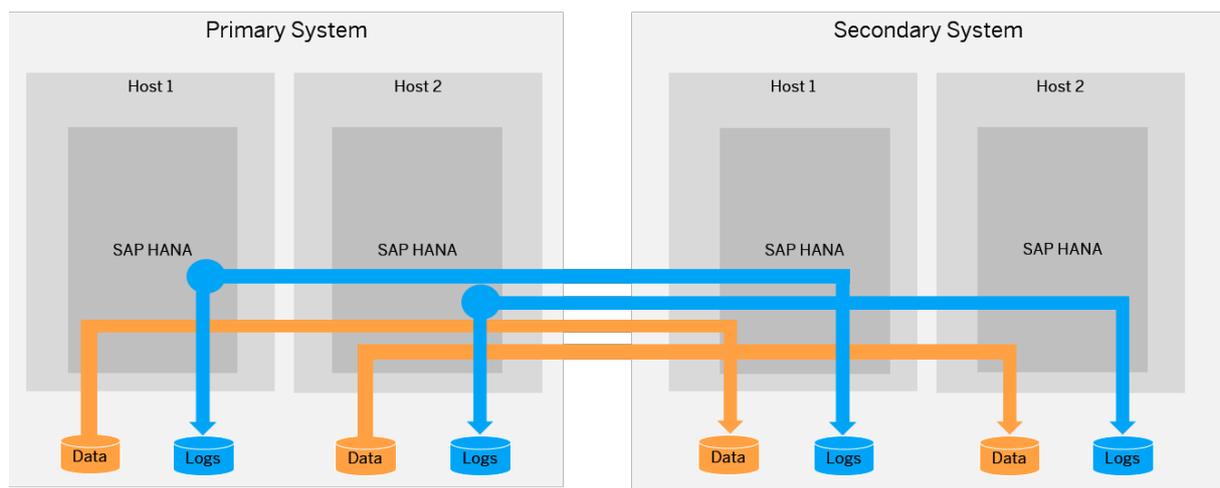
Related Information

[SAP Note 1755396](#)

10.1.1.3 System Replication

System replication is available in every SAP HANA installation offering an inherent disaster recovery support.

System replication is set up so that a secondary system is configured as an exact copy of the active primary system, with the same number of active hosts in each system. The number of standby hosts need not be identical. With multitier system replication you can have a third system attached to the first secondary making it a replication chain of three systems. Each service instance of the primary SAP HANA system communicates with a counterpart in the secondary system. With multitarget system replication the primary system can replicate data changes to more than one secondary system.



The secondary system can be located near the primary system to serve as a rapid failover solution for planned downtime, or to handle storage corruption or other local faults, or, it can be installed in a remote site to be used in a disaster recovery scenario. Also both approaches can be chained together with multitier system replication. Like storage replication, this disaster recovery option requires a reliable connection channel between the primary and secondary sites. The instances in the secondary system operate in recovery mode. In this mode, all secondary system services constantly communicate with their primary counterparts, replicate and persist data and logs, and load data to memory. The main difference to primary systems is that the secondary systems do not accept requests or queries.

When the secondary system is started in recovery mode, each service component establishes a connection with its counterpart, and requests a snapshot of the data in the primary system. From then on, all logged changes in the primary system are replicated. Whenever logs are persisted in the primary system, they are also sent to the secondary system. A transaction in the primary system is not committed until the logs are

replicated. What this means in detail, can be configured by choosing one of the log replication modes. For an overview of the replication modes, see *Replication Modes for SAP HANA*.

If the connection to the secondary system is lost, or the secondary system crashes, the primary system after a brief, configurable, timeout will resume replication. The secondary system persists, but does not immediately replay the received log. To avoid a growing list of logs, incremental data snapshots are transmitted asynchronously from time to time from the primary system to the secondary system. If the secondary system has to take over, only that part of the log needs to be replayed that represents changes that were made after the most recent data snapshot. In addition to snapshots, the primary system also transfers status information regarding which table columns are currently loaded into memory. The secondary system correspondingly preloads these columns. In the event of a failure that justifies full system takeover, an administrator instructs the secondary system to switch from recovery mode to full operation. The secondary system, which already preloaded the same column data as the primary system, becomes the primary system by replaying the last transaction logs, and then starts to accept queries.

i Note

To prevent unauthorized access to the SAP HANA database, the internal communication channels between the primary site and the secondary site in a system replication scenario need to be protected. This may include filtering access to the relevant ports and channels by firewalls, implementing network separation, or applying additional protection at the network level (for example, VPN, IPSec). We recommend routing the connection between the two sites over a special site-to-site high-speed network, which typically already implements security measures such as separation from other network access and encryption or authentication between sites. The details of security measures and implementation of additional network security measures depend on your specific environment. For more information about network and security aspects, see the *SAP HANA Master Guide* and the *SAP HANA Security Guide*.

Related Information

[Configure SAP HANA System Replication: Overview of Steps \[page 1092\]](#)

[Recovery with System Replication \[page 1371\]](#)

[Replication Modes for SAP HANA System Replication \[page 1093\]](#)

10.1.1.4 Service Auto-Restart

Service auto-restart supports fault recovery for one service.

In the event of a software failure or an intentional intervention by an administrator that disables one of the configured SAP HANA services (Index Server, Name Server, and so on), the service will be restarted by the SAP HANA service auto-restart watchdog function, which automatically detects the failure and restarts the stopped service process. Upon restart, the service loads data into memory and resumes its function. While all data remains safe the service recovery takes some time.

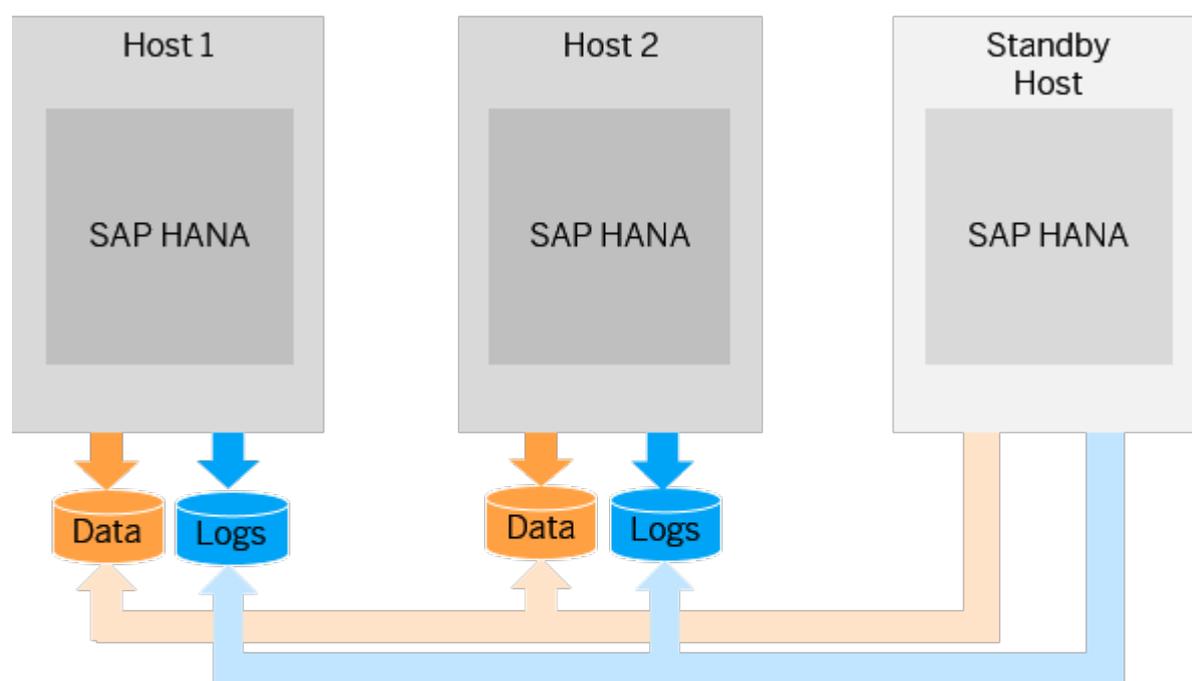
10.1.1.5 Host Auto-Failover

Host auto-failover supports fault recovery for a failed host.

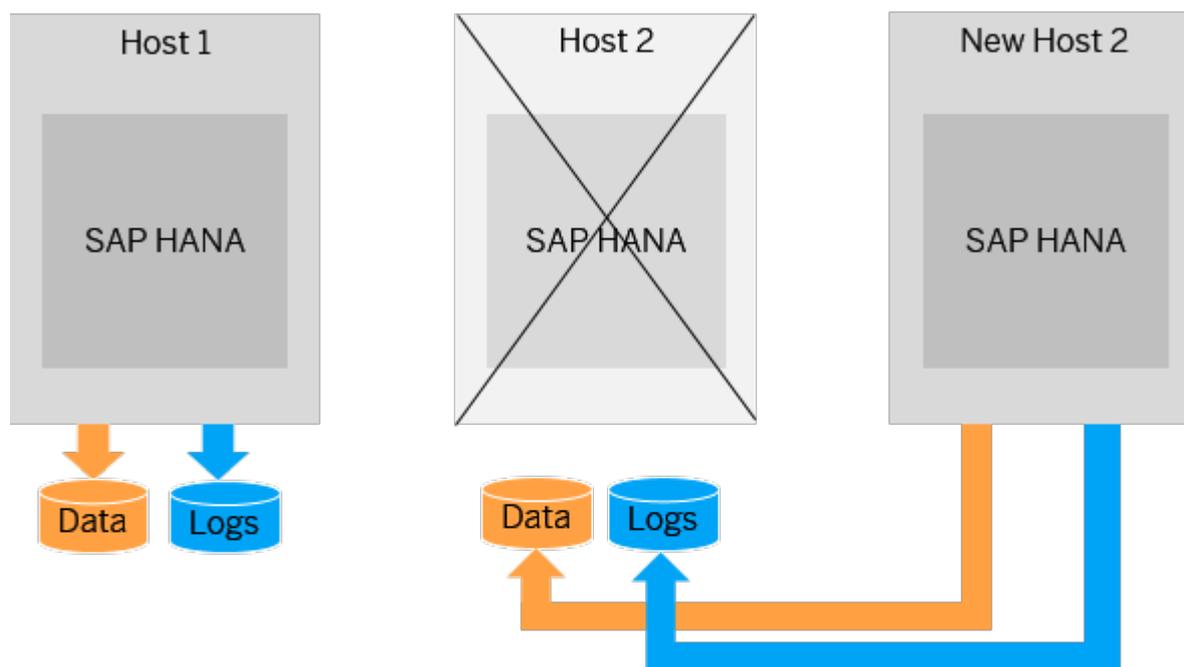
Host auto-failover is a local fault recovery solution that can be used in addition or as an alternative measure to system replication. One (or more) standby hosts are added to an SAP HANA system, and configured to work in standby mode. As long as they are in standby mode the databases on these hosts do not contain any data and do not accept requests or queries. This means they cannot be used for other purposes such as quality or test systems.

When an active (worker) host fails, a standby host automatically takes its place. If neither the name server process `hdbnameserver` nor `hdbdaemon` respond to network requests (because the instance is stopped, the OS has been shut down or powered off), a host is marked as inactive and an auto-failover is triggered. Since the standby host may take over operation from any of the primary hosts, it needs shared access to all the database volumes. This can be accomplished by a shared, networked storage server, by using a distributed file system, or with vendor-specific solutions that use an SAP HANA programmatic interface, the Storage Connector API, to dynamically detach and attach (mount) networked storage upon failover.

This scenario is illustrated in the graphic below:



Once repaired, the failed host can rejoin the system as the new standby host to reestablish the failure recovery capability:



In support of host auto-failover, database clients can be configured with the connection information of multiple hosts, optionally including the standby host. The client connection code (ODBC, JDBC, and so on) will try to connect to one of these, and upon successful connection receives the updated connection configurations. This ensures that clients can continue to reach the SAP HANA database, even after failover.

i Note

In a scale-out system, a failover from the failing master host is triggered automatically even if no standby host is configured. Then a worker host is attached to the volumes of the previous master.

i Note

It is not possible to do a seamless failover. A manual failover to a standby host can be triggered by stopping a worker host using the command `HDB stop`.

Related Information

[Setting Up Host Auto-Failover \[page 1208\]](#)

[SAP HANA - Host Auto-Failover](#)

10.1.2 Configuring SAP HANA System Replication

SAP HANA system replication is a mechanism for ensuring the high availability of your SAP HANA system.

Through the continuous replication of data from a primary to a secondary system, including in-memory loading, system replication facilitates rapid failover in the event of a disaster. Productive operations can be resumed with minimal downtime.

System replication can be set up or managed on the command line with `hdbnsutil`, using the SAP HANA cockpit, the SAP HANA studio, or with SAP Landscape Management (SAP LaMa).

The following administration activities are possible with `hdbnsutil`, using the SAP HANA cockpit, or the SAP HANA studio:

- Performing the initial set-up, that is enabling system replication and establishing the connection between two identical systems
- Monitoring the status of system replication to ensure that both systems are in sync
- Triggering takeover by the secondary system in the event of a disaster and failback once the original system is available again
- Disabling system replication

Note

There is no support for running one site on a little-endian hardware and the other site on a big-endian system in a system replication landscape. In SAP HANA the following endianness is supported in the corresponding SAP HANA and OS versions:

Supported Endianness

SAP HANA Version	Linux OS Version	Endianness
SPS11 & SPS12	Linux Intel SLES 11	little-endian
	Linux Power SLES 11	big-endian
SAP HANA 2.0 SPS00	Linux Intel SLES 12	little-endian
	Linux Power SLES 12	

Furthermore, system replication is supported between Intel little-endian (SAP HANA SPS12 or SAP HANA 2.0 SPS 00) and Power little-endian (SAP HANA 2.0 SPS 00).

Related Information

[SAP Landscape Management \(SAP LaMa\) Documentation](#)

[Orchestrated nZDM for SAP HANA with a single TakeOver](#)

10.1.2.1 General Prerequisites for Configuring SAP HANA System Replication

System replication enables recovery from a data center outage by switching to a secondary site.

The following prerequisites need to be considered:

- The primary and secondary system are both installed and configured. You have verified that both are independently up and running.
- The configuration of hosts in the primary and secondary systems must be the same, that is, the number of hosts must be the same but also the names of the host roles, failover groups and worker groups must be identical in both systems. This implies that if there is a standby host on the primary system it need not be available on the secondary system.
- All configuration steps have to be executed on the master name server node only.
- During an upgrade of the system replication landscape, the software version of the current secondary system has to be equal or newer to the version of the current primary system.

i Note

During a failback, the roles of your system switch. Make sure in this case that your primary system doesn't have a newer software version than the secondary system.

i Note

For active/active (read enabled) setups the SAP HANA versions must be the same on the primary and the secondary system. Use this setup mainly during the upgrade process of the system replication landscape.

- The secondary system must have the same SAP system ID (<SID>) and `instance number` as the primary system.

i Note

The primary replicates all relevant license information to the secondary. An additional license is not required. See SAP Note 2211663.

- System replication between two systems on the same host is not supported.
- The `.ini` file configuration must be similar for both systems. Any changes made manually, or by SQL commands on one system should be manually duplicated on the other system. Automatic configuration parameter checks will alert you to configuration differences between the two systems.

i Note

To keep the `ini` file configuration similar on both systems, the INI parameter checker is per default configured to check for differences. Additionally, it can be configured to replicate parameter changes from the primary system to the secondary system.

- To secure the system replication communication channel between the primary and the secondary system configure the `ini` parameters `[system_replication_communication] / listeninterface` and `allowed_sender` as described in *Host Name Resolution for System Replication*.

- If the host names of the primary and the secondary system are the same (for example, because two systems are used that have identical host names) change the host names used on the secondary system. For more information, see *Rename an SAP HANA System Host*.
- Check that the hostnames in the primary system are different to the hostnames used in the secondary system.
You can see this in the SAP HANA studio, at the end of the environment variable `SAP_RETRIEVAL_PATH` and with the python script `landscapeHostConfiguration.py`
For more information, see *Host Name Resolution for System Replication*
- Ensure that `log_mode` is set to "normal" in the persistence section of the `global.ini` file. Log mode normal means that log segments are backed up.
- You are logged on to both systems as the operating system user (user `<sid>adm`) or you have provided its credentials when prompted.
- You have performed a data backup or storage snapshot on the primary system. In multiple-container systems, the system database and all tenant databases must be backed up. This is necessary to start creating log backups. Activated log backup is a prerequisite to get a common sync point for log shipping between the primary and secondary system.
- Both systems should run on the same endianness platform.
- You must prepare the secondary system for authentication while it is still shut down. To do this, copy the system PKI `SSFS.key` and the `.dat` file from the primary system to the secondary system. For more information, see SAP Note 2369981.

Note

To setup system replication with XS Advanced, copy the following XS advanced secure store files from your primary system to the same location on the secondary system (if already existing, overwrite the corresponding files in the target location):

```
/usr/sap/<SID>/SYS/global/xsa/security/ssfs/data/SSFS_<SID>.DAT
```

```
/usr/sap/<SID>/SYS/global/xsa/security/ssfs/key/SSFS_<SID>.KEY
```

For more information about XS Advanced and system replication, see SAP Note 2300936.

After performing these steps, register the secondary system as secondary and start it.

- Besides updating the XS Advanced keys, you must update also the engine keys:
`/usr/sap/<SID>/SYS/global/security/rsecssfs/data/SSFS_<SID>.DAT`
`/usr/sap/<SID>/SYS/global/security/rsecssfs/key/SSFS_<SID>.KEY`
- SAP HANA dynamic tiering is not supported with multitarget system replication. For more information about SAP HANA system replication with SAP HANA dynamic tiering, see SAP Note 2447994.

Related Information

[Rename an SAP HANA System Host \[page 1034\]](#)

[Host Name Resolution for System Replication \[page 1119\]](#)

[SAP Note 2211663](#)

[SAP Note 2369981](#)

[SAP Note 2300936](#)

[SAP Note 2447994](#)

10.1.2.2 Configure SAP HANA System Replication: Overview of Steps

This topic provides an overview of the steps involved in setting up system replication between two systems, failing over to a secondary system, failing back to a primary system, and disabling system replication.

Procedure

1. Set up system replication on primary and secondary systems:
 - a. Start the primary system.
 - b. Create an initial data backup or storage snapshot on the primary system. In multiple-container systems, the system database and all tenant databases must be backed up.
 - c. Enable system replication on the primary system (sr_enable).
 - d. Prepare the secondary system for authentication by copying the system PKI SSFS .key and the .dat file from the primary system to the secondary system. For more information, see SAP Note 2369981.

To set up system replication with XSA, copy the secure store files from your primary system to the same location on the secondary system before starting the secondary system. For more information, see *General Prerequisites for Configuring System Replication*.
 - e. Register the secondary system with the primary system (sr_register).
 - f. Start the secondary system.
2. During failover, the secondary system takes over from primary system:
 - a. Secondary system in data center B takes over from primary in data center A (sr_takeover).
 - b. Stop primary system in data center A.
 - c. When the primary system is available again, register it with the secondary system (sr_register).

The roles are switched, the original primary is registered as a secondary system. The original secondary is the production system.
 - d. Start the system in data center A.
3. Failback to the original primary system:
 - a. Send a takeover command from the system in data center A (sr_takeover).
 - b. Stop the system in data center B.
 - c. Register the system in data center B as secondary again (sr_register).
 - d. Start the system in data center B.
4. Disable system replication:
 - a. Unregister the secondary system.
 - b. Disable system replication on the primary system.

Related Information

[Rename an SAP HANA System Host \[page 1034\]](#)

[Configuring the Network for Multiple Hosts \[page 1438\]](#)

[Implementing a HA/DR Provider \[page 1215\]](#)

[Enable Data and Log Volume Encryption in an Existing SAP HANA Database \[page 867\]](#)

[SAP Note 2211663](#)

[SAP Note 2369981](#)

10.1.2.3 Replication Modes for SAP HANA System Replication

While registering the secondary system, you need to decide which replication mode to use.

SAP HANA offers different modes for the replication of the redo log:

Replication modes

Log Replication Mode	Description
Synchronous in-memory (SYNCMEM)	<p>The secondary system sends an acknowledgment back to the primary system as soon as the data is received in memory. The disk I/O speed on the secondary system doesn't influence the primary's performance.</p> <p>When the connection to the secondary system is lost, the primary system continues the transaction processing and writes the changes only to the local disk.</p> <p>Data loss can occur when primary and secondary fail at the same time as long as the secondary system is connected or when a takeover is executed, while the secondary system is disconnected. This option provides better performance because it is not necessary to wait for disk I/O on the secondary system, but it is more vulnerable to data loss.</p>
Synchronous (SYNC)	<p>The secondary system sends an acknowledgment back to the primary system as soon as the data is received and persisted to the log volumes on disk.</p> <p>When the connection to the secondary system is lost, the primary system continues the transaction processing and writes the changes only to the local disk. No data loss occurs in this scenario as long as the secondary system is connected. Data loss can occur, when a takeover is executed while the secondary system is disconnected.</p> <p>Additionally, this replication mode can run with a full sync option. This means that log write is successful when the log buffer has been written to the log file of the primary and the secondary systems. When the secondary system is disconnected (for example, because of network failure), the primary system suspends the transaction processing until the connection to the secondary system is reestablished. No data loss occurs in this scenario. You can set the full sync option for system replication with the parameter <code>[system_replication]/enable_full_sync</code>. For more information on how to enable the full sync option, see <i>Enable Full Sync Option for System Replication</i>.</p>
Asynchronous (ASync)	<p>The primary system sends redo log buffers to the secondary system asynchronously. The primary system commits a transaction when it has been written to the log file of the primary system and sent to the secondary system through the network. It doesn't wait for confirmation from the secondary system.</p> <p>This option provides better performance because it is not necessary to wait for log I/O on the secondary system. Database consistency across all services on the secondary system is guaranteed. However, it is more vulnerable to data loss. Data changes may be lost on takeover.</p>

i Note

If you plan to add SAP HANA dynamic tiering to your landscape in the future, please check supported replication modes in *SAP Note 2447994* before you enable SAP HANA system replication.

The replication mode can be changed without going through a full data shipping from the primary system to the secondary system afterwards.

To change the replication mode use the following command on the online or offline secondary system:

```
hdbnsutil -sr_changemode --mode=sync|syncmem|async
```

In the `M_SERVICE_REPLICATION` view you can check with the following command whether the replication mode was changed correctly:

```
hdbnsutil -sr_state --sapcontrol=1
```

Related Information

[Full Sync Option for SAP HANA System Replication \[page 1108\]](#)

10.1.2.4 Operation Modes for SAP HANA System Replication

While registering the secondary system, you need to decide in which operation mode to run SAP HANA system replication.

System replication can be run in three operation modes: `delta_datashipping`, `logreplay` or `logreplay_readaccess`. Depending on the configured operation mode, the database sends different types of data packages to the secondary system. For more information, see *Data Transferred to the Secondary System*.

The default operation mode is `logreplay`. With the operation modes `logreplay` and `logreplay_readaccess` no delta data shippings are necessary anymore, the takeover time has been reduced, and more components are initialized at replication time.

In a multitier or multitarget system replication it is not possible to combine `logreplay` and `delta_datashipping` operation modes. In a multitarget system replication only `logreplay` and maximum one `logreplay_readaccess` are supported.

Overview of Operation Modes

Operation Mode	Description
<code>delta_datashipping</code>	This mode establishes a system replication where occasionally (per default every 10 minutes) a delta data shipping takes place in addition to the continuous log shipping. The shipped redo log is not replayed on the secondary site. During takeover the redo log needs to be replayed up to the last arrived delta data shipment.

Operation Mode	Description
logreplay	In this operation mode a redo log shipping is done after system replication was initially configured with one full data shipping. The redo log is continuously replayed on the secondary system immediately after arrival making this step superfluous during takeover. This mode does not require delta data shippings. Because of this, the amount of data which needs to be transferred to the secondary system is reduced.
logreplay_readaccess	This mode is required for replication to an Active/Active (read enabled) secondary system. It is similar to the logreplay operation mode regarding the continuous log shipping, the redo log replay on the secondary system, as well as the required initial full data shipping and the takeover.

Logreplay

Before you begin preparing a replication strategy for an SAP HANA system, you should be aware of the following important aspects regarding the operation modes `logreplay` and `logreplay_readaccess`.

- Registering a secondary with operation mode `logreplay` against a primary running on an SAP HANA revision less than or equal to SPS10 will not work, because the primary does not yet support this feature. Furthermore, for operation mode `logreplay_readaccess` the primary must be running on a revision SAP HANA 2 SPS00 or higher.
- In a NZDU (Near Zero Downtime Upgrade) from an SAP HANA revision less than or equal to SPS 10 to SPS 11 when registering the original primary (failback) after upgrade of the secondary only the operation mode `delta_datashipping` will work, because the former primary's version does not yet support logreplay.
- With the logreplay operation modes if the connection to the secondary is not available, the primary system will keep writing the redo log segments in the online log area to be prepared for the delta log shipping after the connection is reestablished. These log segments are marked as *RetainedFree* until the secondary is in sync again. In this case there is a risk that the log volume may run full. To prevent this:
 - If a secondary is not used anymore, it must be unregistered (`sr_unregister`).
 - If a takeover to the secondary was done, the former primary should be disabled (`sr_disable`).

See *How to Avoid Log Full Situations* in the *SAP HANA Troubleshooting and Performance Analysis Guide* for more details.
- The logreplay operation modes do not support history tables.

i Note

If you plan to add SAP HANA dynamic tiering to your landscape in the future, please check supported operation modes in *SAP Note 2447994* before you enable SAP HANA system replication.

Switching Operation Modes

You can switch operation modes using the `hdbnsutil -sr_register` command and explicitly setting the new operation mode with the `-operationMode` option:

```
hdbnsutil -sr_register --name=<secondary_alias>
```

```
--remoteHost=<primary_host> --remoteInstance=<primary_systemnr>  
--replicationMode=[sync|syncmem|async]--operationMode=[delta_datashipping|  
logreplay|logreplay_readaccess]
```

To change the operation mode; the secondary must be offline. When switching to logreplay or logreplay_readaccess operation modes no full data shipping is necessary. Full data shipping is necessary, however, when switching from logreplay back to delta_datashipping.

Related Information

[Data Transferred to the Secondary System \[page 1096\]](#)

[Active/Active \(Read Enabled\) \[page 1157\]](#)

[SAP HANA System Replication Command Line Reference \[page 1117\]](#)

10.1.2.4.1 Data Transferred to the Secondary System

Depending on the configured operation mode, the database sends different types of data packages to the secondary system.

When system replication is configured, the following types of data packages can be sent to the secondary system:

- Initial full data shipping
A full set of data created as an SAP HANA in-place snapshot on the disk of the primary system is initially sent when system replication is set up.
- Delta data shipping
The data that has changed since the last full or the last delta data shipping is transported from time to time from the data area of the primary system to the data area of the secondary system. The default time is every 10 minutes.
When using logreplay and logreplay_readaccess delta data shippings are not required.
- Continuous redo log shipping
Every committing write transaction on the primary system generates redo log buffers, which are continuously sent to the secondary system.

i Note

With the ini file parameter `datashipping_parallel_channels` (default 4) the full and the delta data shipping are done using parallel network channels. You can change it on the secondary system in the `global.ini` section `[system replication]`.

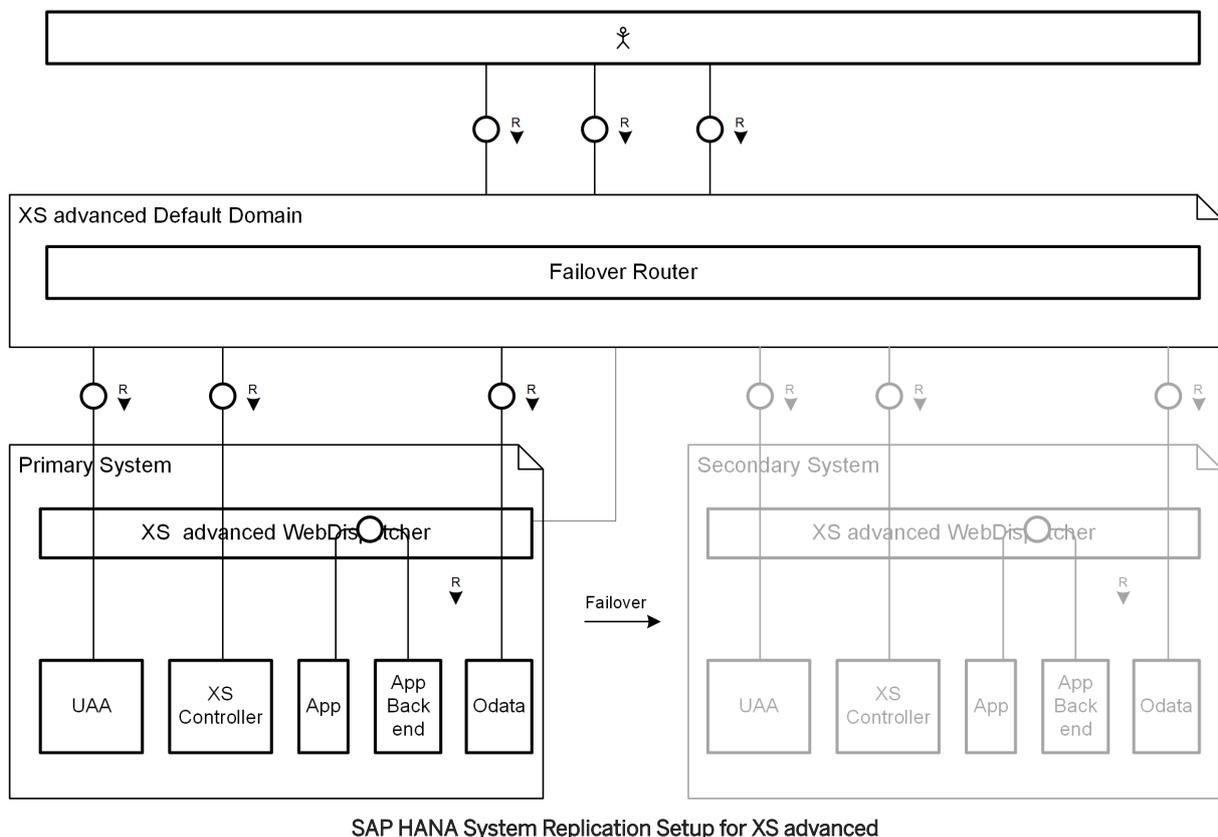
10.1.2.5 SAP HANA System Replication Setup for XS Advanced Runtime

In a system replication setup, all the data – including XS advanced runtime system data and application data – is replicated to a secondary system.

XS advanced services and applications run only on the currently active system. On the secondary system XS advanced services are in an idle state until the takeover takes place.

After the takeover all XS advanced services are started which in turn brings up all XS advanced applications on the secondary system. Moreover, XS advanced services and applications will use the same domains and certificates that were present in the primary system before the takeover started.

For this to work, the central point for XS advanced requests must be the same on the primary and the secondary systems. This is established by using a failover router similar to the high availability setup. For more information on the high availability setup, see *SAP HANA High Availability Setup for XS Advanced Runtime*.



SAP HANA System Replication Setup for XS advanced

In case the failover router terminates SSL, the same rules apply as described in *SAP HANA High Availability Setup for XS Advanced Runtime*.

For more information, see SAP Note 2300936.

Related Information

[Host Auto-Failover Setup with XS Advanced Runtime \[page 1213\]](#)

10.1.2.6 Setting Up SAP HANA System Replication

System replication enables recovery from a data center outage by switching to a secondary site.

You can configure system replication using the following tools:

- SAP HANA cockpit
For more information, see *Set Up System Replication with the SAP HANA Cockpit*.
- SAP HANA studio
For more information, see *Set Up System Replication with the SAP HANA Studio*.
- hdbnsutil
For more information, see *Set Up System Replication with hdbnsutil*.

Related Information

[Set Up SAP HANA System Replication from the Primary System \[page 1099\]](#)

[Set Up SAP HANA System Replication with the SAP HANA Studio \[page 1102\]](#)

[Set Up SAP HANA System Replication with hdbnsutil \[page 1104\]](#)

[Full Sync Option for SAP HANA System Replication \[page 1108\]](#)

10.1.2.6.1 Set Up SAP HANA System Replication with the SAP HANA Cockpit

To set up SAP HANA system replication in the SAP HANA cockpit, first enable system replication on the primary system and then register the secondary system.

Prerequisites

- You have considered all the general prerequisites needed to set up system replication. For more information, see *General Prerequisites for Setting Up SAP HANA System Replication*.
- You have added both systems in the SAP HANA cockpit.

Context

The [System Replication](#) tile on the system overview page provides the possibility to configure system replication. Once the configuration is done, the tile displays information on the operation mode, the replication mode, the configuration type, and the status of system replication.

The secondary system can be registered from the primary system or from the [Overview](#) page of the SAP HANA cockpit. You can register again a previously stopped secondary system when a full data shipping is needed or when you want to change the operation mode

Related Information

[General Prerequisites for Configuring SAP HANA System Replication \[page 1090\]](#)

[Set Up SAP HANA System Replication from the Primary System \[page 1099\]](#)

[Set Up SAP HANA System Replication from the Primary and the Secondary Systems \[page 1100\]](#)

[Reinitialize the Secondary System \[page 1102\]](#)

10.1.2.6.1.1 Set Up SAP HANA System Replication from the Primary System

To set up SAP HANA system replication, first enable system replication on the primary system and then register the secondary system. Use the SAP HANA cockpit to execute these separate configuration steps from the primary system.

Context

This topic describes how to enable system replication on the primary system and then register the secondary system from the primary system.

i Note

If you plan to add SAP HANA dynamic tiering to your landscape in the future, please see [SAP Note 2447994](#) before you enable HANA system replication. SAP HANA dynamic tiering requires certain communication ports, operation modes, and replication modes.

Procedure

1. On the [Overview](#) page of the primary system, choose the [System Replication](#) tile. Initially this tile displays the message *System replication is not yet enabled for this system*.

The *System Replication* page opens. If you performed a data backup before enabling system replication, this page displays overview information on the primary system on the top left and the *Configure System Replication* link on the top right.

2. Choose *Configure System Replication*.

The *System Replication Configuration* dialog opens, allowing you to run the configuration in background.

3. Enter the logical name used to represent the primary system in the *Tier 1 System Details* screen area.
4. Enter the logical name used to represent the secondary system in the *Tier 2 System Details* screen area.
5. Select the secondary system host and mark the checkbox below this area to stop the system.
6. Select a replication mode. For more information on the available replication modes, see *Replication Modes for SAP HANA System Replication*.
7. Select an operation mode. For more information on the available operation modes, see *Operation Modes for SAP HANA System Replication*.
8. Decide whether to initiate data shipping or not.
9. Check *Start Secondary after Registration*.
10. Optional: To add a new system to your system replication landscape configuration click *Add Tier 3 System* on the bottom left.
11. Choose *Configure System Replication*.

The *System Replication Configuration* dialog opens, allowing you to run the configuration in background.

Related Information

[General Prerequisites for Configuring SAP HANA System Replication \[page 1090\]](#)

[Replication Modes for SAP HANA System Replication \[page 1093\]](#)

[Operation Modes for SAP HANA System Replication \[page 1094\]](#)

[SAP Note 2447994](#)

10.1.2.6.1.2 Set Up SAP HANA System Replication from the Primary and the Secondary Systems

To set up SAP HANA system replication, first enable system replication on the primary system and then register the secondary system. Use the SAP HANA cockpit to execute these configuration steps on the primary system and separately on the secondary system.

Context

This topic describes how to enable system replication on the primary system and then register the secondary system using the SAP HANA cockpit.

Procedure

1. On the *Overview* page of the primary system, choose the *System Replication* tile. Initially this tile displays the message *System replication is not yet enabled for this system*.

The *System Replication* page opens. If you performed a data backup before enabling system replication, this page displays overview information on the primary system on the top left and the *Configure System Replication* link on the top right.

2. Enter the logical name used to represent the primary system and choose *Configure* on the bottom right.

The *System Replication Configuration* dialog opens, allowing you to run the configuration in background.

3. On the *Overview* of the future secondary system, choose the *Overall Database Status* tile.
4. Choose *Stop System* on the bottom right, because the system has to be offline in order to be registered as a secondary system.

Back on the *Overview* of the future secondary system the *Overall Database Status* tile displays the status *Stopped*.

5. On the *Overview* page of the secondary system, choose the *System Replication* tile.

The *System Replication* page opens, displaying overview information on the secondary system on the top left and the *Register Secondary System* button on the top right.

6. Choose *Register Secondary System*.

The *System Replication Configuration* page opens.

7. On the *System Replication Configuration* page enter the logical name used to represent the secondary system.
8. On the *System Replication Configuration* page select a replication mode. For more information on the available replication modes, see *Replication Modes for SAP HANA System Replication*.
9. Select an operation mode. For more information on the available operation modes, see *Operation Modes for SAP HANA System Replication*.
10. Enter the host of the source system.

i Note

If you are operating a distributed system on multiple hosts, enter the name of the host on which the master name server is running.

11. Check *Start Secondary after Registration*.
12. Review the configured information and choose *Configure* on the bottom right.

The *System Replication Configuration* dialog opens. After the configuration is complete, the *System Replication Overview* page displays information on the configured systems.

Related Information

[General Prerequisites for Configuring SAP HANA System Replication \[page 1090\]](#)

[Replication Modes for SAP HANA System Replication \[page 1093\]](#)

[Operation Modes for SAP HANA System Replication \[page 1094\]](#)

10.1.2.6.1.3 Reinitialize the Secondary System

You can register again a previously stopped secondary system using the SAP HANA cockpit.

Context

You can register again a previously stopped secondary system. You must do this when a full data shipping is needed or when you want to change the operation mode.

Procedure

1. On the *Overview* page of the stopped secondary system, choose the *System Replication* tile.
2. On the *System Replication Overview*, choose *Reinitialize Secondary System* on the top right.
3. On the *System Replication Configuration* page, you can now change the configuration. Change the operation mode or resync the persistencies using the *Initiate full data shipping* option.

The secondary system is up and running again.

10.1.2.6.2 Set Up SAP HANA System Replication with the SAP HANA Studio

To set up SAP HANA system replication between two identical SAP HANA systems, you must first enable system replication on the primary system and then register the secondary system.

Prerequisites

- You have considered all the general prerequisites needed to set up system replication. For more information, see *General Prerequisites for Setting Up SAP HANA System Replication*.
- You have added both systems in the SAP HANA studio.

Procedure

1. Enable system replication on the primary system, which has to be online, as follows:
 - a. In the *Systems* view, right-click the primary system and choose ► *Configuration and Monitoring* ► *Configure System Replication* >.

The *Configure System Replication* dialog opens. The *Enable System Replication* option is selected by default.

i Note

You can also access the *Configure System Replication* dialog from the ► *Landscape* ► *System Replication* ► tab.

- b. Choose *Next*.
 - c. Enter the logical name used to represent the primary system and choose *Next*.
 - d. Review the configured information and choose *Finish*.
 - e. Stop the secondary with right-click on the secondary system and choosing ► *Configuration and Monitoring* ► *Stop System* ►.
2. Register the secondary system as follows:
- a. Stop the secondary system if it is still running. Right-click the secondary system and choose ► *Configuration and Monitoring* ► *Stop System* ►
 - b. In the *Systems* view, right-click the secondary system and choose ► *Configuration and Monitoring* ► *Configure System Replication* ►.
The *Configure System Replication* dialog opens.
 - c. Choose *Register Secondary System* and then *Next*.
 - d. Enter the required system information and the logical name used to represent the secondary system.

i Note

If you are operating a distributed system on multiple hosts, you enter the name of the host on which the master name server is running.

- e. Specify the log replication mode. For more information on the available replication modes, see *Replication Modes for SAP HANA System Replication*.
 - f. Specify the operation mode. For more information on the available operation modes, see *Operation Modes for SAP HANA System Replication*.
 - g. Review the configured information and choose *Finish*.
3. Optional: Configure the parameters in the `system_replication` section of the `global.ini` file.
These parameters determine for example the size and frequency of data and log shipping requests. All parameters have a default configuration.
4. If necessary, start the secondary system.

i Note

The secondary system is started automatically unless you deselected the corresponding option during configuration (step 2).

The secondary system requests an initial full data replica from the primary system.

Results

You have enabled system replication and registered the secondary system with the primary system. The secondary system operates in recovery mode. All secondary system services constantly communicate with

their primary counterparts, replicate and persist data and logs, and load data to memory. However, the secondary system does not accept SQL connections.

In the *Systems* view, the primary and secondary systems appear as operational (■). If the secondary system is not open for read access, it appears as operational (■) but with an error (✖) indicating that no connection to the database is available. For more information, see *Generic Conditions for Active/Active (Read Enabled)*.

Related Information

[General Prerequisites for Configuring SAP HANA System Replication \[page 1090\]](#)

[Replication Modes for SAP HANA System Replication \[page 1093\]](#)

[Operation Modes for SAP HANA System Replication \[page 1094\]](#)

[Add an SAP HANA System \[page 122\]](#)

[Stop a System \[page 183\]](#)

[SAP HANA System Replication Configuration Parameters \[page 1109\]](#)

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 1317\]](#)

[Rename an SAP HANA System Host \[page 1034\]](#)

[Enable Data and Log Volume Encryption in an Existing SAP HANA Database \[page 867\]](#)

[Generic Conditions for Active/Active \(Read Enabled\) \[page 1159\]](#)

[SAP Note 611361](#)

10.1.2.6.3 Set Up SAP HANA System Replication with hdbnsutil

You can configure SAP HANA system replication with `hdbnsutil`.

Prerequisites

You have considered all the general prerequisites needed to set up system replication. For more information, see *General Prerequisites for Setting Up SAP HANA System Replication*.

Procedure

1. Enable system replication on the primary system as follows:
 - a. In the Administration editor of SAP HANA studio, choose the *Configuration* tab and ensure that `log_mode` is set to "normal" in the `persistence` section of the `global.ini` file.

Log mode `normal` means that log segments must be backed up. Log mode `overwrite` means log segments are freed by the savepoint (therefore only useful for test installations without backup and recovery).

- b. Do an initial data backup or create a storage snapshot. In multiple-container systems, the system database and all tenant databases must be backed up.
- c. As `<sid>adm` on the command line enable the primary for system replication and give it a logical name with the following command. The primary system must be online at this time:

```
cd /usr/sap/<sid>/HDB<instancenr>/exe
```

```
./hdbnsutil -sr_enable --name=<primary_alias>
```

Option Name	Value	Description
<code>--name</code>	<code><primary_alias></code>	Alias used to represent your primary site and assign it as the primary site for system replication

To check if the site has been successfully enabled for system replication with `hdbnsutil` run:

```
cd /usr/sap/<sid>/HDB<instancenr>/exe
```

```
./hdbnsutil -sr_state
```

- d. Stop the secondary system:

```
sapcontrol -nr <instance_number> -function StopSystem HDB
```

2. Register the secondary system as follows:

- a. Enable system replication on the secondary system as user `<sid>adm` with the following command:

```
hdbnsutil -sr_register --name=<secondary_alias>
--remoteHost=<primary_host> --remoteInstance=<primary_systemnr>
--replicationMode=[sync|syncmem|async] --operationMode=[delta_datashipping|
logreplay|logreplay_readaccess]
```

`hdbnsutil -sr_register` Call Options

Option Name	Value	Description
<code>--name</code>	<code><secondary_alias></code>	Alias used to represent the secondary site
<code>--remoteHost</code>	<code><primary_host></code>	Name of the primary host that the secondary registers with
<code>--remoteInstance</code>	<code><primary_instancenr></code>	Instance number of primary
<code>--replicationMode</code>	<code>[sync syncmem async]</code>	Log replication modes
<code>--operationMode</code>	<code>[delta_datashipping logreplay logreplay_readaccess]</code>	Log operation mode

To check if the site has been successfully enabled for system replication with hdbnsutil run:

```
cd /usr/sap/<sid>/HDB<instancenr>/exe
```

```
./hdbnsutil -sr_state
```

- b. Start the secondary system to reinitialize it with the following command:

As <sid>adm:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function  
StartSystem HDB
```

Related Information

[General Prerequisites for Configuring SAP HANA System Replication \[page 1090\]](#)

[Rename an SAP HANA System Host \[page 1034\]](#)

[Host Name Resolution for System Replication \[page 1119\]](#)

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 1317\]](#)

10.1.2.6.4 Initializing the Secondary

After the secondary system has been registered with the primary site it is initialized with the data from the primary site.

There are two general situations that can occur:

- The secondary site is completely unrelated to the primary site
- The secondary is related to the primary site as it was:
 - Already registered before as secondary and was shut down for a time.
 - A former primary site, in this case, the system is prepared for failback by replicating in the opposite direction.

If the secondary system is unrelated to the primary system, a full data shipping is done. An in-place snapshot created on the disk of the primary system is initially sent to the secondary system. This initial full data shipping can be prevented by manual intervention and the secondary system can be initialized with a binary storage copy of the primary system's persistence. For more information see, *Initialize the Secondary with Storage Copy from Primary*.

If the persistence (that is data and log volumes) of the secondary site is related to the primary site (it actually contains the persistence of the primary at a former time), the newly registered site can be synced with a delta data or log shipping.

After a new registration of the secondary site, a delta data or log shipping is always attempted. For more information, see *Resync Optimization*.

Related Information

[Initialize the Secondary with Storage Copy from Primary \[page 1107\]](#)

[Resync Optimization \[page 1144\]](#)

10.1.2.6.4.1 Initialize the Secondary with Storage Copy from Primary

The secondary site can be initialized using a binary storage copy from the primary site.

Context

For this procedure copy only the data, not the log.

Procedure

1. Create a consistent binary storage copy from the primary system for the persistence of all services. You can use the snapshot technology to create an IO consistent persistence copy. Create a full copy of the persistence using the IO consistent storage snapshots.

If you cannot use the method above, create a consistent OS copy of persistence while the primary system is stopped.

2. Shut down the secondary system.
3. Transfer or mount the full copy on the secondary system.
4. Replace the persistence of the secondary site by the storage copy from the primary site.
5. Register the secondary system without [--force_full_replica].
6. Start the secondary system.

Results

When the secondary system is started after the new registration, the initialization optimizations are carried out. The system checks if the persistence of the secondary site is compatible with the persistence of the primary site. The secondary system checks if its persistence is compatible with the persistence of the primary site. If this check succeeds the secondary system requests only a delta data shipping.

10.1.2.6.5 Full Sync Option for SAP HANA System Replication

When activated, the full sync option for SAP HANA system replication ensures that a log buffer is shipped to the secondary system before a commit takes place on the local primary system.

The full sync option can be enabled for SYNC replication (that is not for SYNCMEM). With the activated full sync option, transaction processing on the primary blocks when the secondary is currently not connected and newly created log buffers cannot be shipped to the secondary site. This behavior ensures that no transaction can be locally committed without shipping the log buffers to the secondary site. The full sync option can be switched on and off using the command:

```
hdbnsutil -sr_fullsync --enable|--disable
```

It changes the setting of the parameter `enable_full_sync` in the `system_replication` section of the `global.ini` file accordingly. However, in a running system, full sync does not become active immediately. This is done to prevent the system from blocking transactions immediately when setting the parameter to true. Instead, full sync has to first be enabled by the administrator. In a second step it is internally activated, when the secondary is connected and becomes ACTIVE.

In the system view `M_SERVICE_REPLICATION` the setting of the full sync option can be viewed in the column "FULL_SYNC" using SQL.

It can have the following values:

- **DISABLED:** Full sync is not configured at all. The parameter `enable_full_sync = false` in the `system_replication` section of the `global.ini` file.
- **ENABLED:** Full sync is configured, but it is not yet active, so transactions do not block in this state. To become active the secondary has to connect and `REPLICATION_STATUS` has to be ACTIVE.
- **ACTIVE:** Full sync mode is configured and active. If the network connection to a connected secondary is closed, transactions on the primary side will block in this state.

If full sync is enabled when an active secondary is currently connected, the `FULL_SYNC` will be immediately set to ACTIVE.

If the secondary is stopped, disable `FULL_SYNC`. Otherwise the primary blocks and it is not possible to stop it.

i Note

Resolving a blocking situation of the primary caused by the enabled full sync option must be done with the `hdbnsutil` command, since a configuration changing command could also block in this state. This is also necessary, if you want to shut down the currently blocking primary. Otherwise it is not possible to stop it.

In a multitarget system replication setup, configure the full sync option on the primary system. Enter the site name used for the secondary system when registering it:

```
global.ini
[system_replication]
enable_full_sync[<secondary_site_name>] = true
```

i Note

In a multitarget system replication setup, you can use `hdbnsutil -sr_fullsync` only for turning off the full sync option.

Related Information

[SAP HANA System Replication Command Line Reference \[page 1117\]](#)

10.1.2.6.6 SAP HANA System Replication Configuration Parameters

Several configuration parameters are available for configuring SAP HANA system replication between the primary and secondary system, for example, the size and frequency of data and log shipping requests.

The system replication parameters are defined in the `system_replication` section of the `global.ini` file and have the default values shown below. The *System* column defines whether the parameter can be set on the primary, the secondary, or both.

Note

`preload_column_tables` uses the Boolean keywords "true" or "false". Numbers do not work in place of the keywords.

Parameter	Type	Unit	Default	System	Description
<code>datashipping_min_time_interval</code>	int	seconds	600 (10 min)	Secondary	Minimum time interval between two data shipping requests from secondary system. If <code>datashipping_logsize_threshold</code> (see next parameter) is reached first, the data shipping request will be sent before the time interval is elapsed, when the log size threshold is reached.
<code>datashipping_logsize_threshold</code>	int	bytes	5*1024*1024*1024 (5GB)	Secondary	Minimum amount of log shipped between two data shipping requests from secondary system. If the time defined by <code>datashipping_min_time_interval</code> (see previous parameter) has passed before reaching this threshold, the data shipping request will be sent before this threshold is reached, when the time interval has elapsed.
<code>preload_column_tables</code>	bool	(true/false)	true	Primary and secondary	If set preload of column table main parts is activated. If set on the primary system, the loaded table information is collected and stored in the snapshot that is shipped. If set on the secondary system, this information is evaluated and the tables are actually preloaded there according to the information received on the loaded tables.

Parameter	Type	Unit	Default	System	Description
<code>datashipping_snapshot_max_retention_time</code>	int	minutes	300	Primary	<p>Maximum retention time (in minutes) of the last snapshot that has been completely shipped to the secondary system. Shipped snapshots older than <code>datashipping_snapshot_max_retention_time</code> will be dropped automatically. Snapshots currently used in data shipping are not affected and are not dropped, if data shipping takes longer than <code>datashipping_snapshot_max_retention_time</code>. They can be dropped if data shipping has been finished. If the parameter is set to 0, snapshots are immediately dropped after data replication finishes.</p> <p>When roles are switched between primary and secondary sites in prepare for a fail back later on, the secondary can be initialized with a delta replica between this snapshot and the current persistent state on the "new primary" after takeover. In order to do this:</p> <ul style="list-style-type: none"> • A snapshot has to exist on the new secondary when it starts up for the first time as secondary • The snapshot has to be compatible with the persistence of the new primary. <p>It is verified, if the snapshot has been the source of the primary system before takeover. It cannot be used, if the secondary is registered with an incompatible primary system. If both conditions are true, the secondary can be initialized with a delta replica.</p>

Parameter	Type	Unit	Default	System	Description
logshipping_timeout	int	seconds	30	Primary	<p>Number of seconds, the primary waits for the acknowledgment after sending a log buffer to the secondary site. If the primary does not receive the acknowledgment for a sent log buffer within the time defined by <code>logshipping_timeout</code>, it will close the connection to the secondary site in order to continue data processing. This is done to prevent the primary system from blocking transaction processing if there is a hang situation on the connection to the secondary site. After the timeout period for a send operation has elapsed transactions are written only on primary side until the secondary has reconnected. The <code>logshipping_timeout</code> does not define a blocking period for logshipping on the primary site in general. It is used to close hanging connections on the primary site, that are not getting automatically closed. If there is a connection close from the secondary site detected, transaction processing will immediately continue without waiting for the timeout to be elapsed. This can happen any time, also when the primary is currently not waiting for acknowledges from the secondary site. If the primary site should block in all situations, when the connection to the secondary site is getting lost, the full sync option should be used. In this case the primary system will stop</p>
logshipping_async_buffer_size	int	bytes	67108864 (64MB)	Primary	<p>In asynchronous replication mode, the log writer copies the log buffers first into an intermediate memory buffer and continues processing. A separate thread reads log buffers from this memory buffer and sends them to the secondary site asynchronously over the network.</p> <p>This parameter determines, how much log can be intermediately buffered. This buffer is especially useful in peak times, when log is generated faster than they can be sent to the secondary site. If the buffer is large, it can handle peaks for a longer time period.</p> <p>The behavior of buffer full situations can be controlled by the parameter <code>logshipping_async_wait_on_buffer_full</code></p> <p>The parameter can be changed online, but will become active the next time the secondary system reconnects.</p>

Parameter	Type	Unit	Default	System	Description
logshipping_asy nc_wait_on_buff er_full	bool	true/ false	true	Primary	<p>This parameter controls the behavior of the primary/source system in asynchronous log shipping mode, when the log shipping buffer is full.</p> <p>If set to true, the primary/source system potentially waits, until there is enough space in the log shipping buffer, so that the log buffer can be copied into it. This can slow down the primary system, if there is currently high load that cannot be handled by the network connection.</p> <p>If the parameter is set to false, the connection to the secondary system will be closed temporarily in order not to impact the primary system. Later, the secondary can reconnect and sync using delta shipping.</p>
reconnect_time_ interval	int	seconds	30	Secondary	<p>If a secondary system is disconnected from the primary system due to network problems, the secondary tries to reconnect periodically after the time interval specified in this parameter has passed.</p>
enable_full_syn c	bool	bool	false	Primary	<p>If set, system replication operates in full sync mode when the replication mode SYNC is set. In full sync mode, transaction processing blocks, when the secondary is currently not connected and newly created log buffers cannot be shipped to the secondary site. This behavior ensures that no transaction can be locally committed without shipping to the secondary site.</p>

Parameter	Type	Unit	Default	System	Description
enable_log_compression	bool	true/false	false	Secondary	<p>If activated, log buffers will be compressed before sending them over the network to the secondary site. The secondary site decompresses the log buffers it receives and then writes them to disk. If network bandwidth is the bottleneck in the system replication setup log buffer compression can improve log shipping performance because less data is being sent over the network.</p> <p>The drawback to sending a compressed log buffer to the secondary site is that it requires additional time and processing power for compression and decompression. This can result in worse log shipping performance if turned on in a configuration with a fast network.</p> <p>The parameter has to be set on the secondary site. It can be changed online, but the secondary system has to re-connect to the primary site in order to activate the parameter change.</p>
enable_data_compression	bool	true/false	false	Secondary	<p>If activated, data pages will be compressed before sending them over the network to the secondary site. The secondary site decompresses the data pages it receives and then writes them to disk. If network bandwidth is the bottleneck in the system replication setup data compression can improve log shipping performance because less data is being sent over the network.</p> <p>The drawback to sending compressed data pages to the secondary site is that it requires additional time and processing power for compression and decompression. This can result in worse data shipping performance if turned on in a configuration with a fast network.</p> <p>The parameter has to be set on the secondary site. It can be changed online, but the secondary system has to re-connect to the primary site in order to activate the parameter change.</p>

Parameter	Type	Unit	Default	System	Description
keep_old_style_alert	bool	true/ false	true	Primary	Before SPS 09 closed replication connections and configuration parameter mismatches were alerted with Alert 21. With SPS 09 two dedicated alerts have been introduced for both error situations for better monitoring. By default old style alerting is still offered for backwards compatibility. When setting this parameter to false, the old behavior is turned off and only new alerts will be generated.
operation_mode	enum		logreplay	Secondary	<p>Operation mode of the secondary site during replication. There are three different settings for this parameter:</p> <ul style="list-style-type: none"> • delta_datashipping System Replication uses data and log shipping for replication. Log buffers received by the secondary site are just saved to disk, savepoints after intermediate delta data shippings truncate the log. Column table merges are not executed on the secondary site, but merged tables on the primary site are transported via delta data shippings to the secondary site. This operation mode is available since SPS 05. • logreplay System Replication uses an initial data shipping to initialize the secondary site. After that only log shipping is done and log buffers received by the secondary are replayed there. Savepoints are executed individually for each service and column table merges are executed on the secondary site. • logreplay_readaccess System Replication uses an initial data shipping to initialize the secondary site. After that only log shipping is done and log buffers received by the secondary are replayed there. Savepoints are executed individually for each service and column table merges are executed on the secondary site. Furthermore, read only access via SQL is possible to the secondary system.

Parameter	Type	Unit	Default	System	Description
enable_log_retention	enum		auto	Primary, Secondary	<p>Enables/Disables log retention on a system replication site. Log retention on the primary site is useful when the secondary should sync with the primary by re-shipping missing log after a network outage or downtime. If the missing log is not available anymore on the primary site a data shipping is required (delta in operation mode delta_datashipping, full in all other operation modes). Log retention on the secondary site is needed to keep log for optimized re-sync during failback.</p> <p>There are three configuration options:</p> <ul style="list-style-type: none"> • auto Log retention is automatically enabled, if the secondary is in operation mode logreplay or logreplay_readaccess, it is disabled by default for operation mode delta_datashipping. • on Log retention is enabled • off Log retention is disabled <p>When log retention is enabled and the system is configured as system replication primary site, then the primary will not free log segments when the secondary site is disconnected. When setting log retention explicitly to on/off it should also be set for operation mode delta_datashipping or for failback with delta log shipping optimization. In the latter case after takeover to the secondary the old primary can re-sync via missing log with the new primary site and no full data shipping is required for initialization.</p> <p>In a multitarget system replication configuration, if <code>enable_log_retention = force_on_takeover</code> is configured, the log will be retained during replication for all direct secondaries until a takeover is executed. During takeover, the parameter is set to force. This means the log will be retained independently of any secondary system. For more information, see <i>Log Retention and Multi-target System Replication</i> in <i>Log Retention</i>.</p>

Parameter	Type	Unit	Default	System	Description
logshipping_max_retention_size	int	MB	1048576 (1TB)	Primary	<p>Set the maximum amount of log that will be kept for syncing a system replication secondary system. This value only has an effect, if log retention is enabled.</p> <p>Two situations have to be distinguished here:</p> <p>If logshipping_max_retention_size has been set to a value other than 0, when no secondary is connected log segments are not reused even if they are truncated and backed up until the max size limit has been reached or the system runs into a log full situation. If the max size limit is reached or in log full situation segments that are only kept for syncing the secondary site will be reused. This setting prevents the system from hanging on the primary site due to too many log segments, that are held for syncing the secondary site. With this setting the primary is kept running with the drawback that the secondary cannot sync anymore.</p> <p>If logshipping_max_retention_size is configured to 0, then log segments required for secondary syncing are not reused and a log full results in a system standstill on primary site until log writing can continue. This setting allows you to assign a higher priority to being able to sync the secondary over a standstill on the primary. When the reason for the log full has been resolved (on primary or secondary site), transaction processing can continue.</p> <div data-bbox="884 1317 1396 1576" style="background-color: #f0f0f0; padding: 10px;"> <p>i Note</p> <p>The default setting logshipping_max_retention_size = 1048576 (MB) of 1 TB means that 1 TB of size is configured for every service, which replicates data to a secondary system (that is, every service owning a persistence in form of data and log volume).</p> </div> <div data-bbox="884 1592 1396 1901" style="background-color: #f0f0f0; padding: 10px;"> <p>❖ Example</p> <p>If the services nameserver, two indexeservers (for example, two tenant databases) and an xsengine are running in your SAP HANA system, the total configured log retention size will be 4 TB (4 x 1 TB). With this setting it can happen that the disk full is reached before the Retained-Free marked log segments are overwritten.</p> </div>

Parameter	Type	Unit	Default	System	Description
					<p>If you want to change the default value of 1 TB, you can do this in the global.ini. Another option is to set this parameter in the service ini files individually. For example, if the value is set in the global.ini of the SystemDB, in the global.ini of a tenant database, and in the indexserver.ini of a tenant database, the indexserver.ini setting would win and will be taken for log retention of this indexserver.</p>
datashipping_parallel_channels	int		4	Secondary	<p>The parameter defines the number of network channels used by full or delta datashipping. The actual number of channels for each shipping can be adjusted by the system to reduce overhead depending on the current amount of data to be shipped.</p> <p>Higher parallelism can be useful when large amount of data (above several GB at least) needs to be shipped, and the utilization of network bandwidth by single network stream is low. Please note that the overall bandwidth is still limited by the I/O bandwidth, because the data needs to be read from the primary system.</p> <p>To deactivate the parameter, change the default to 0.</p>

Related Information

[Change a System Property in SAP HANA Studio \[page 301\]](#)
[Log Retention \[page 1145\]](#)

10.1.2.6.7 SAP HANA System Replication Command Line Reference

This topic provides an overview of SAP HANA system replication commands and options.

sr_commands

Command	Options	System	Online/Offline	Description
-sr_enable	[--name=<site alias>]	Primary	Online	<p>Enables a site to serve as a system replication source site.</p> <p>In multitier setups the --name= option is mandatory on the second tier. If you register the tier two as the source system for the tier three system do not use this option with -sr_enable as you have already done this as part of -sr_register.</p>
-sr_disable		Primary	Online	Disables system replication capabilities on source site.
-sr_register	--remoteHost=<primary master host> --remoteInstance=<primary instance id> --replicationMode=sync syncmem async --operationMode=delta_datashipping logreplay logreplay_readaccess --name=<unique site name> [--force_full_replica]	Secondary	Offline	Registers a site to a source site and creates the replication path for the system replication. Specifies the replication mode. Specifies the operation mode. Specify the site name. If parameter is given, a full data shipping is initiated. Otherwise a delta data shipping is attempted.

Command	Options	System	Online/Offline	Description
-sr_unregister	[--id=<site id> --name=<site name>]	Primary	Secondary offline, Primary online (to remove metadata)	<p>Unregisters a secondary site from its source.</p> <p>You can use this command to unregister the secondary from its source from the secondary system.</p> <p>Using the options for site id and site name you can unregister the secondary by executing the command on the primary system.</p>
-sr_changemode	--mode=sync syncmem async	Secondary	Online and offline	Changes the replication mode of a secondary site.
-sr_takeover		Secondary	Online and offline	Switches system replication primary site to the calling site.
-sr_state		Primary and Secondary	Online and offline	Shows status information about system replication site.
-sr_cleanup		Primary	Offline	Removes system replication configuration.

Related Information

[SAP Note 1945676: Correct usage of hdbnsutil -sr_unregister](#) 

10.1.2.6.8 Host Name Resolution for System Replication

The correct mapping of internal host names between primary and secondary systems is required for system replication.

With SAP HANA system replication, each SAP HANA instance communicates on the service level with a corresponding peer in the secondary system to persist the same data and logs as in the primary system. The replication of the transactional load can be configured to work in synchronous or asynchronous mode, depending mainly on the distance between the two sites. For a full description of system replication, see the section on high availability and the white paper *Introduction to High Availability for SAP HANA*.

Communication between the primary and the secondary system is based on internal host names. The host names of the other site must always be resolvable, either through configuration in SAP HANA or corresponding entries in the `/etc/hosts` file.

For system replication it is not necessary to edit the `/etc/hosts` file, internal ('virtual') host names must be mapped to IP addresses in the `global.ini` file to create a dedicated network for system replication; the syntax for this is as follows:

```
global.ini
[system_replication_hostname_resolution]
<ip-address_site>=<internal-host-name_site>
<...>
```

This is necessary to ensure that each site can resolve the host name of other replicating sites and that hosts can be switched seamlessly in the event of a takeover. These virtual host names must be set before registering the secondary system because the `-sr_register` command uses this mapping.

These entries in the `[system_replication_hostname_resolution]` section are used in combination with the `listeninterface` parameter in the `[system_replication_communication]` section which in a replication scenario can be set to either `.global` or `.internal`. There is another `listeninterface` parameter in the `[communication]` section which is required for the communication between SAP HANA services (name server, index server, and so on) in a distributed system, but it has no impact on system replication (see following illustrations).

Virtual host names can also be used if the hostnames have a domain suffix, for example: the internal hostnames can be defined as `ab820*` and `ab830*`, but the public names have to include the domain, such as `ab820*.abc.xyz.com` and `ab830*.def.xyz.com`

The following table shows two settings of the `listeninterface` parameter and the corresponding mapped host names. These are both illustrated in the following graphics showing a multi-node replication environment with a separate internal network for replication.

⚠ Caution

Note that by default no mappings are specified and system replication communication uses the default network route (typically the public network). If you use a public network instead of a separate network, you **must** secure this connection with additional measures such as a firewall or a virtual private network and/or TLS/SSL.

listeninterface	Host Name Resolution Mappings	Additional Information
<code>.global</code>	IP addresses and host names of neighboring sites (minimum) or for all hosts of own site as well as for all hosts of neighboring sites. Also applies to multi-tier setups.	This establishes a separate network for system replication communication. → Tip Also for multitier and multitarget setups, this is how you can use a dedicated network for system replication communication.

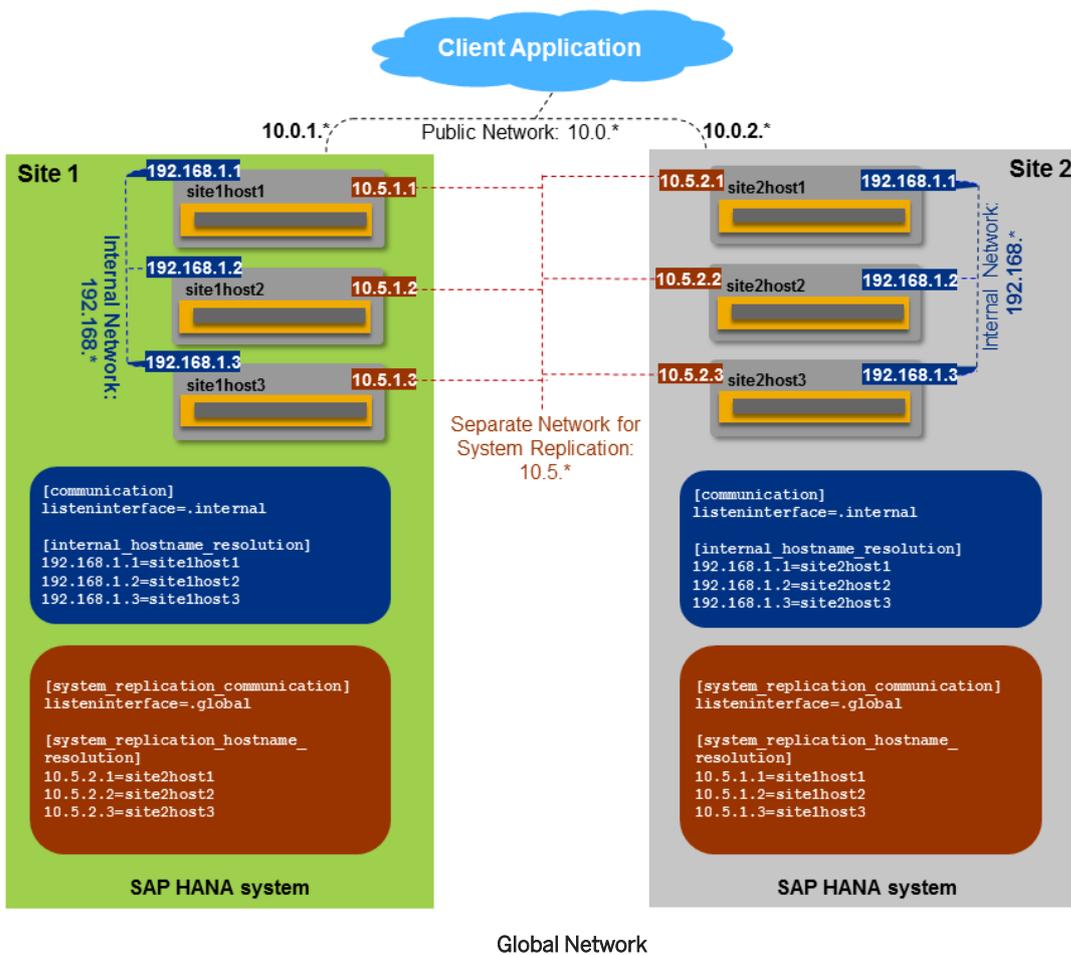
listeninterface	Host Name Resolution Mappings	Additional Information
<code>.internal</code>	Entries for all hosts of own site as well as for all hosts of neighboring sites	<p>A separate network is used for system replication communication. The primary hosts listen on the dedicated ports of the separate network only, and incoming requests on the public interfaces are rejected.</p> <div style="border: 1px solid orange; padding: 5px;"> <p>⚠ Caution</p> <p>As of SAP HANA 1.0 SPS 11, network communication for system replication with <code>listeninterface=.internal</code> is supported for two-tier replication but not for multitier and multitar-get setups.</p> </div>

Examples

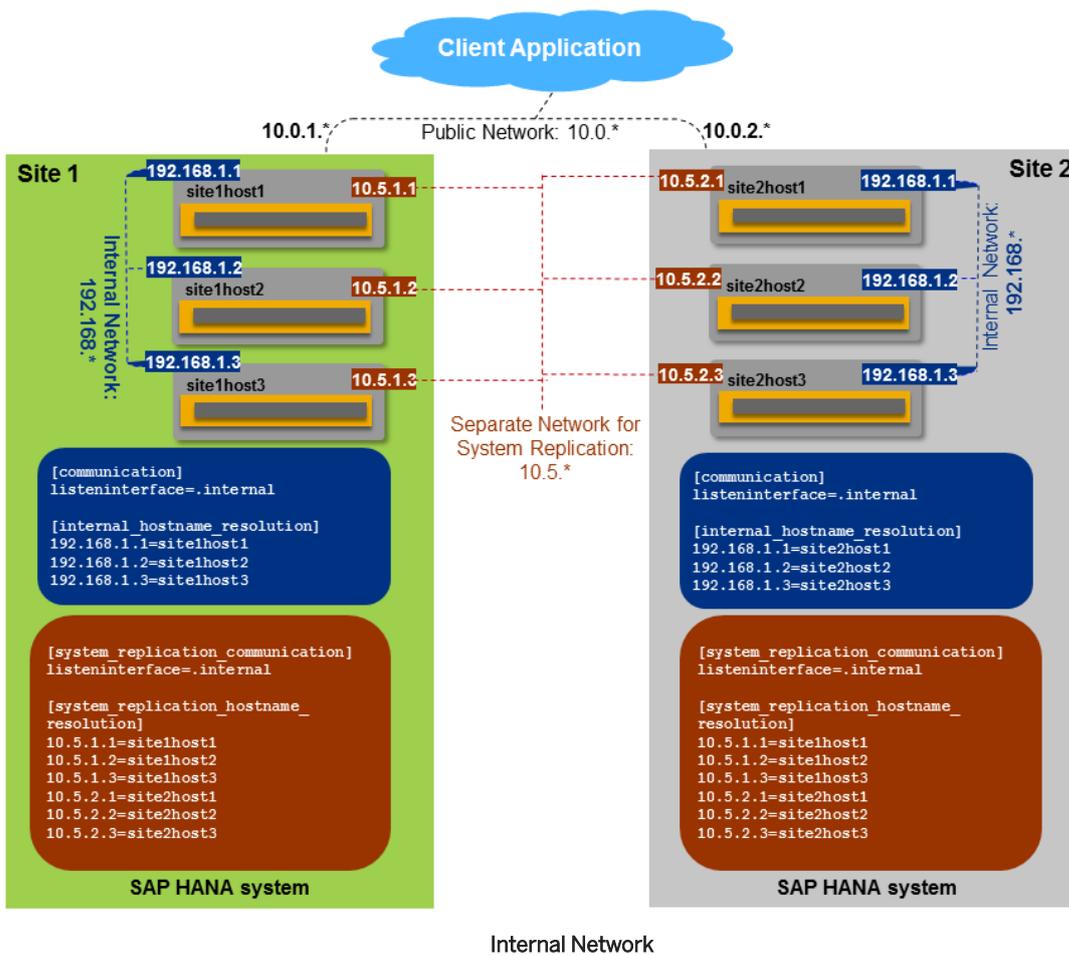
The following examples show the host name resolution configuration for system replication to a secondary site. Three distinct networks can be identified:

- Public network with addresses in the range of 10.0.1.*
- Network for internal SAP HANA communication between hosts at each site: 192.168.1.*
- Dedicated network for system replication: 10.5.1.*

In the first example, the `[system_replication_communication]listeninterface` parameter has been set to `.global` and only the hosts of the neighboring replicating site are specified.



In the following example, the `[system_replication_communication]listeninterface` parameter has been set to `.internal` and all hosts (of both sites) are specified.



Important Security Information

If you use a public network instead of a separate internal network and/or the parameter `listeninterface=.global`, you **must** secure this connection with additional measures such as a firewall or a virtual private network and/or TLS/SSL.

If no separate internal network channel is configured for SAP HANA system replication, the `allowed_sender` parameter can be used to restrict communication between the primary and secondary sites to certain hosts. For this purpose, the following settings are made in the `global.ini` file on the primary site:

```
global.ini
[system_replication_communication]
allowed_sender=<list of IP-addresses of secondary or CIDR-netmasks>
```

An example of this parameter value would be `10.0.1.0/30`. The default is no restriction.

For more security-related information, see the *SAP HANA Security Guide* and *SAP HANA Checklists and Recommendations*.

Related Information

[System Replication \[page 1085\]](#)

[Set Up SAP HANA System Replication with hdbnsutil \[page 1104\]](#)

[SAP Note 2036111: Configuration parameters for SAP HANA \(including system replication\)](#)

[White paper "Introduction to High Availability for SAP HANA"](#)

[How-To Guide: How To Configure Network Settings for HANA System Replication](#)

10.1.2.6.9 Data and Log Compression

SAP HANA system replication supports a number of compression methods for log and data shipping.

The following types of compression for log and data shipping are supported:

- Log
 - Log buffer tail compression
 - Log buffer content compression
- Data
 - Data page compression

Log Buffer Tail Compression

All log buffers are aligned to 4kb boundaries by a filler entry. With log buffer tail compression the filler entry is cut off from the buffer before sending it over the network and added again when the buffer has reached the secondary site. So only the net buffer size is transferred to the secondary site.

The size of the filler entry is less than 4kb, this is the maximum size reduction per sent log buffer. If the size of the log buffers is quite large, the compression ratio is quite limited. Log buffer tail compression is turned on by default.

Log Buffer and Page Content Compression

As of SPS 09 log buffers and data pages shipped to the secondary site can be compressed using a lossless compression algorithm (lz4). By default content compression is turned off. You can turn it on by setting the following configuration parameters on the secondary site in the system_replication section of the global.ini file:

- `enable_log_compression = true`
- `enable_data_compression = true`

Log and data compression is especially useful when system replication is used over long distances, for example using the replication mode ASYNC.

The open source compression algorithm lz4 has been selected because of its speed and compression ratios and the time overhead introduced for compression/decompression is quite low. Log buffer content

compression works also in combination with log buffer tail compression. So only the content part of the log buffer is compressed, without considering the filler entry.

Related Information

[External link to LZ4](#) ➔

10.1.2.6.10 Example Set Up of SAP HANA System Replication

This example shows you how to set up system replication with a single host system.

Context

To set up system replication with two hosts you may have to change the hostnames.

In this example a single host system is used, in multi-host systems all hosts have to be renamed.

i Note

To rename hosts in a production system replication landscape, system replication must be first deactivated. This means you have to first unregister and disable the secondary system before renaming any hosts. Once you have renamed the hosts then you can enable recovery mode again and register the secondary system with the primary system to re-activate system replication.

Procedure

1. Enable system replication on the primary system, with the hostname ej11.

```
cd /usr/sap/<sid>/HDB<instancenr>/exe
```

```
./hdbnsutil -sr_enable --name=dcsite1
```

2. Stop the secondary system. The primary system can stay online.

As <sid>adm

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StopSystem  
HDB
```

3. Register the secondary system with the following command:

```
cd /usr/sap/<sid>/HDB<instancenr>/exe
```

```
./hdbnsutil -sr_register
```

```
--name=dcsite2
--remoteHost=ej11
--remoteInstance=50
--mode=sync
```

Also see *SAP Note 611361 Hostnames of SAP servers*

4. Start the secondary system. This initiates the initial data transfer.

As `<sid>adm`

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StartSystem
HDB
```

Related Information

[Rename an SAP HANA System Host \[page 1034\]](#)

[SAP Note 611361](#)

10.1.2.7 Performing a Takeover

During a takeover you switch your active system from the current primary system to the secondary system.

If your primary data center is not available, due to a disaster or for planned downtime for example, and a decision has been made to fail over to the secondary data center, you can perform a takeover on your secondary system.

We recommend that you should use third-party external tools to be able to check if hosts, the network, and data center are still available.

In addition, you can use python scripts to decide when a takeover should be carried out. For a detailed description of the available python scripts, see *Takeover Decision Based on SAP HANA Python Scripts*.

To help you decide if a takeover is advisable, see the decision tree in *SAP Note 2063657*.

You can perform a takeover using the following tools:

- SAP HANA cockpit
For more information, see *Perform a Takeover with the SAP HANA Cockpit*.
- SAP HANA studio
For more information, see *Perform a Takeover with the SAP HANA Studio*.
- hdbnsutil
For more information, see *Perform a Takeover with hdbnsutil*.

Related Information

[SAP Note 2063657](#)

[Takeover Decision Based on SAP HANA Python Scripts \[page 1127\]](#)

[Perform a Takeover with the SAP HANA Cockpit \[page 1130\]](#)

[Perform a Takeover with the SAP HANA Studio \[page 1130\]](#)

[Perform a Takeover with hdbnsutil \[page 1131\]](#)

10.1.2.7.1 Takeover Decision Based on SAP HANA Python Scripts

You can use python scripts to decide when a takeover should be carried out.

landscapeHostConfiguration.py

The landscapeHostConfiguration.py script shows the status of the primary system:

- SAP HANA is OK
- SAP HANA will be OK after a host auto-failover, for example
- Or not enough instances are started and a takeover would be useful

Note

The script does not tell you if the secondary system is ready for a takeover.

The script provides the following tabular output. It also provides an overall status and a return code to match the overall host status.

The return codes of the script are:

Return code	Description
0	Fatal Internal script error, the state could not be determined
1	Error
2	Warning
3	Info
4	OK

A takeover is only recommended when the return code from the script is 1 (error).

Overall host status: OK.

```
h04adm@ld8520:/usr/sap/H04/HDB04/exe/python_support> python landscapeHostConfiguration.py
| Host | Host | Host | Failover | Remove | Storage | Failover | Failover | NameServer | NameServer | IndexServer | IndexServer |
| | Active | Status | Status | Status | Partition | Config Group | Actual Group | Config Role | Actual Role | Config Role | Actual Role | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| ld8520 | yes | ok | | | | 1 | default | default | master 1 | master | worker | master |
| ld8521 | yes | ok | | | | 2 | default | default | master 2 | slave | worker | slave |
| ld8522 | yes | ignore | | | | 0 | default | default | master 3 | slave | standby | standby |
overall host status: ok
```

Overall host status: Warning. This is because a Host Auto-Failover is taking place.

```
h04adm@1d8520:/usr/sap/H04/HDB04/exe/python_support> python landscapeHostConfiguration.py
| Host | Host | Host | Failover | Remove | Storage | Failover | Failover | NameServer | NameServer | IndexServer | IndexServer | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Active | Status | Status | Status | Status | Partition | Config Group | Actual Group | Config Role | Actual Role | Config Role | Actual Role |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 1d8520 | yes | ok | | | | 1 | default | default | master 1 | master | worker | master |
| 1d8521 | no | warning | failover to 1d8522 | | | 2 | default | default | master 2 | slave | worker | slave |
| 1d8522 | yes | ignore | | | | 0 | default | default | master 3 | slave | standby | standby |
overall host status: warning
```

Overall host status: Information. The landscape is completely functional, but the actual role of the host differs from the configured role.

```
h04adm@1d8520:/usr/sap/H04/HDB04/exe/python_support> python landscapeHostConfiguration.py
| Host | Host | Host | Failover | Remove | Storage | Failover | Failover | NameServer | NameServer | IndexServer | IndexServer | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Active | Status | Status | Status | Status | Partition | Config Group | Actual Group | Config Role | Actual Role | Config Role | Actual Role |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 1d8520 | yes | ok | | | | 1 | default | default | master 1 | master | worker | master |
| 1d8521 | no | info | | | | 0 | default | default | master 2 | slave | worker | standby |
| 1d8522 | yes | info | | | | 2 | default | default | master 3 | slave | standby | slave |
overall host status: info
```

Overall host status: Error. There are not enough active hosts.

```
h04adm@1d8520:/usr/sap/H04/HDB04/exe/python_support> python landscapeHostConfiguration.py
| Host | Host | Host | Failover | Remove | Storage | Failover | Failover | NameServer | NameServer | IndexServer | IndexServer | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Active | Status | Status | Status | Status | Partition | Config Group | Actual Group | Config Role | Actual Role | Config Role | Actual Role |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 1d8520 | yes | ok | | | | 1 | default | default | master 1 | master | worker | master |
| 1d8521 | no | error | | | | 2 | default | default | master 2 | slave | worker | slave |
| 1d8522 | yes | ignore | | | | 0 | default | default | master 3 | slave | standby | standby |
overall host status: error
```

Note

In the event of a network split, a so called "split brain" scenario, the script cannot tell if the instance in the other half of the network is fully functional. Therefore an automatic takeover decision should not be based on this script alone.

`systemReplicationStatus.py`

The `systemReplicationStatus.py` script shows the status of system replication.

Using `systemReplicationStatus.py` has the advantage of showing whether the secondary systems are in sync or not. This provides more confidence if a takeover is justified because if system replication was never in sync or is outdated, unexpected loss of data might occur.

However, rather than calling both scripts manually and calculate the required action based on return codes, you can use the `getTakeoverRecommendation.py` script.

getTakeoverRecommendation.py

This script has three return codes:

landscapeHostConfiguration	systemReplicationStatus	Takeover Status	Reason
Error/Fatal/Warning	NoHsr/Error/Unknown/Initializing/Syncing/Active	Required	Primary system has errors
OK/Info/Ignore	NoHsr/Error/Unknown/Initializing/Syncing	Cannot decide	Unknown system replication status
OK/Info/Ignore	Active	Possible	Primary system is up and system replication is in sync

When the `getTakeoverRecommendation` script is called, it shows the takeover recommendation based on the current system state. However, when the primary system faces any error situation, the system replication status cannot be determined anymore. Thus, the previous state should be saved and compared against the current state.

❁ Example

This is a sample implementation of a python script that uses `getTakeoverRecommendation` to act as a minimalistic cluster manager:

```
import time
import subprocess
from getTakeoverRecommendation import TakeoverDecision
def main():
    wasInSync = False
    while True:
        recommendation =
        subprocess.call(["python", "getTakeoverRecommendation.py", "--sapcontrol=1"])
        if not wasInSync and recommendation is TakeoverDecision.Required:
            print "Primary defect & no sync => NO TAKEOVER"
        if wasInSync and recommendation is TakeoverDecision.Required:
            print "Primary defect & sync => TAKEOVER"
        nowInSync = recommendation is TakeoverDecision.Possible
        wasInSync = nowInSync
```

The output depends on the previous state with the result of the current call of `getTakeoverRecommendation`. If no sync state is reached, a takeover is not advised. But once the systems are in sync, the next error of the primary system will suggest a takeover. Any subsequent negative return value will reset the sync state as it is no longer ensured that the replicated data is current.

10.1.2.7.2 Perform a Takeover with the SAP HANA Cockpit

You can perform a takeover on your secondary system using the SAP HANA cockpit.

Prerequisites

You have the credentials of the operating system user (<sid>adm user) that was created when the system was installed.

Context

The *System Replication* tile provides information on the operation mode, the configuration type, and the status of system replication.

Procedure

1. On the *Overview* page of the secondary system meant to perform the takeover, choose the *System Replication* tile.

The *System Replication* tile opens displaying the *System Replication Overview*.

2. Choose *Take Over* on the top right.
3. To start the takeover, click *Ok* in the *System takes over* dialog.

10.1.2.7.3 Perform a Takeover with the SAP HANA Studio

You can perform a takeover on your secondary system using the SAP HANA studio.

Prerequisites

You are logged on to the secondary system as the operating system user (user <sid>adm) or can enter these credentials when prompted.

Procedure

1. In the *Systems* view, right-click the secondary system and choose ► *Configuration and Monitoring* ► *Configure System Replication* ►.
2. Choose *Perform Takeover*.
3. Enter the required system information and choose *Next*.
4. Review the information and choose *Finish*.
5. If necessary, stop the primary system.

i Note

If the primary system is still running at the time of takeover, it is stopped automatically unless you deselected the corresponding option during takeover (step 3).

Results

The secondary system is now the production system. If the system is already running, it comes out of recovery mode and becomes fully operational immediately: it replays the last transaction logs and starts to accept queries. If the system is offline, it takes over production operation when you start it.

10.1.2.7.4 Perform a Takeover with hdbnsutil

You can perform a takeover on your secondary system with the hdbnsutil.

Context

The takeover command can be executed both when the secondary system is in an offline state or online state. The secondary site must be fully initialized. You can check this in M_SERVICE_REPLICATION or in SAP HANA studio ► *Administration Console* ► *Landscape* ► *System Replication* ►. The secondary site is ready for takeover if all services display *REPLICATION_STATUS ACTIVE*.

If the secondary system is online, it can be brought out of recovery mode and become fully operational as follows:

Procedure

As `<sid>adm` enter the following command to enable the secondary system to take over and become the primary system:

```
cd /usr/sap/<sid>/HDB<instancenr>/exe
```

```
./hdbnsutil -sr_takeover
```

If the system is offline, the takeover is actually carried out when the system is next started.

Next Steps

i Note

If you are performing a takeover as part of a planned downtime you should first make sure that the primary system has been fully stopped before performing a takeover to the secondary system.

Related Information

[Stop a System \[page 183\]](#)

[Monitoring SAP HANA Systems During Stop and Start \[page 187\]](#)

10.1.2.7.5 Client Connection Recovery after Takeover

To allow for continued client communication with the SAP HANA system your high availability solution has to also support client connection recovery. Connection recovery after disaster recovery can be done with network-based IP redirection or network-based DNS redirection.

After a takeover the new primary database server is not aware of previous connections which existed between clients and the former primary server. If the client application does not issue a new request and keeps waiting for a reply from the server, it will not receive an explicit request to close these connections from either of the servers and will keep waiting indefinitely. To prevent this, the SAP HANA client library supports the TCP keepalive feature provided by the operating system. This feature will lead the client to abort the invalid connection on its end and to trigger a reconnect after a specified period during which the former primary server is not reachable.

However, the default keepalive settings for the operating system (2 hours) may lead the client processes to wait for a long time before they abort the connection on their end and trigger a reconnect with the new primary. For

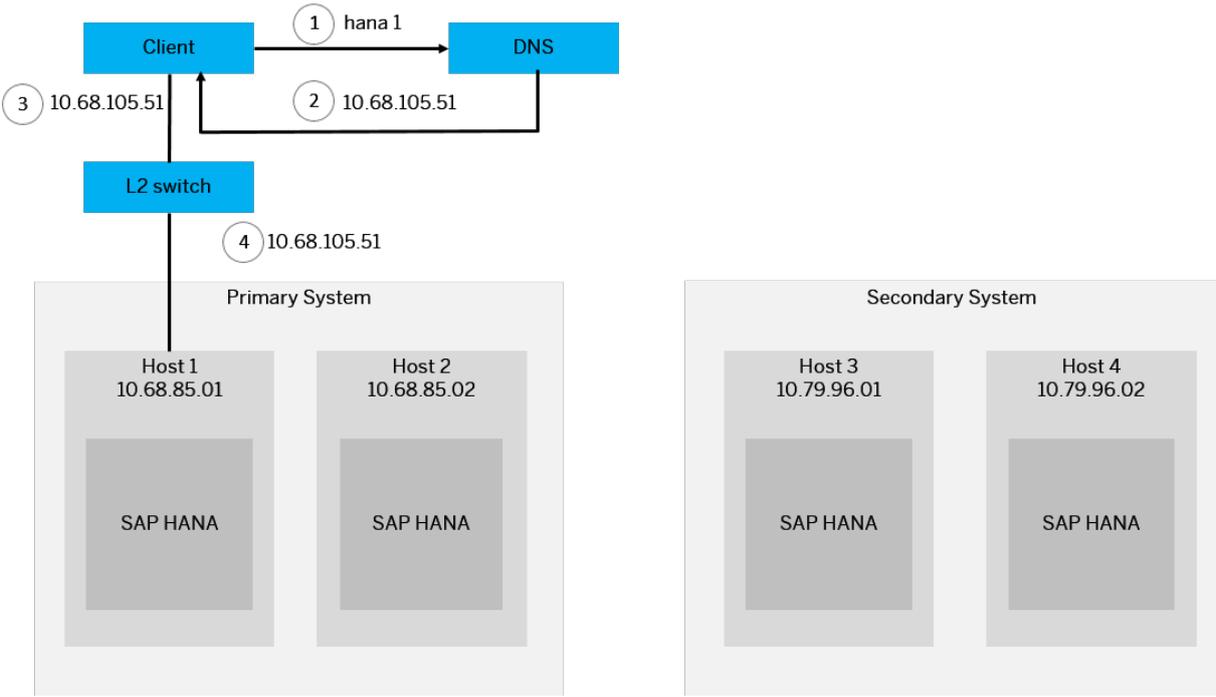
example, the default Linux settings leave the clients waiting for more than two hours before aborting the connection. For more information, see SAP Note 2053504. For instructions on how to configure the keepalive settings to match your needs, see the corresponding documentation for your operating system.

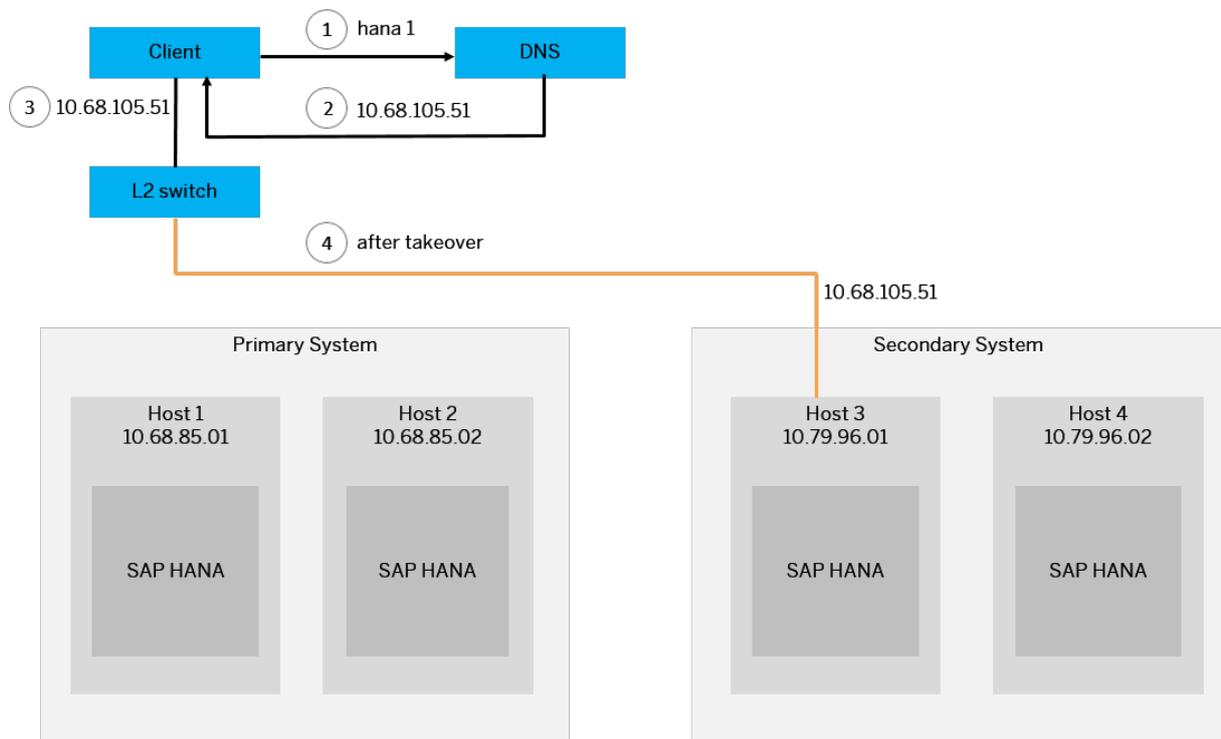
As part of disaster recovery planning you need to consider how IP addresses used by the clients accessing your systems can be moved between primary and secondary systems.

There are different possibilities for enabling client connection recovery.

Network-based IP Redirection

The principle of IP redirection is to define an additional "logical" host name (hana1, in the diagram below) with its own separate logical IP address (10.68.105.51), and then map this initially to the MAC address of the original host in the primary system (by binding it to one of the host's interfaces). As part of the takeover procedure, a script is executed which re-maps the unchanged logical IP address to the corresponding takeover host in the secondary system. This must be done pair-wise, for each host in the primary system. The remapping affects the L2 (OSI layer 2: data link) switching, as can be seen in step 4 of the following diagram:





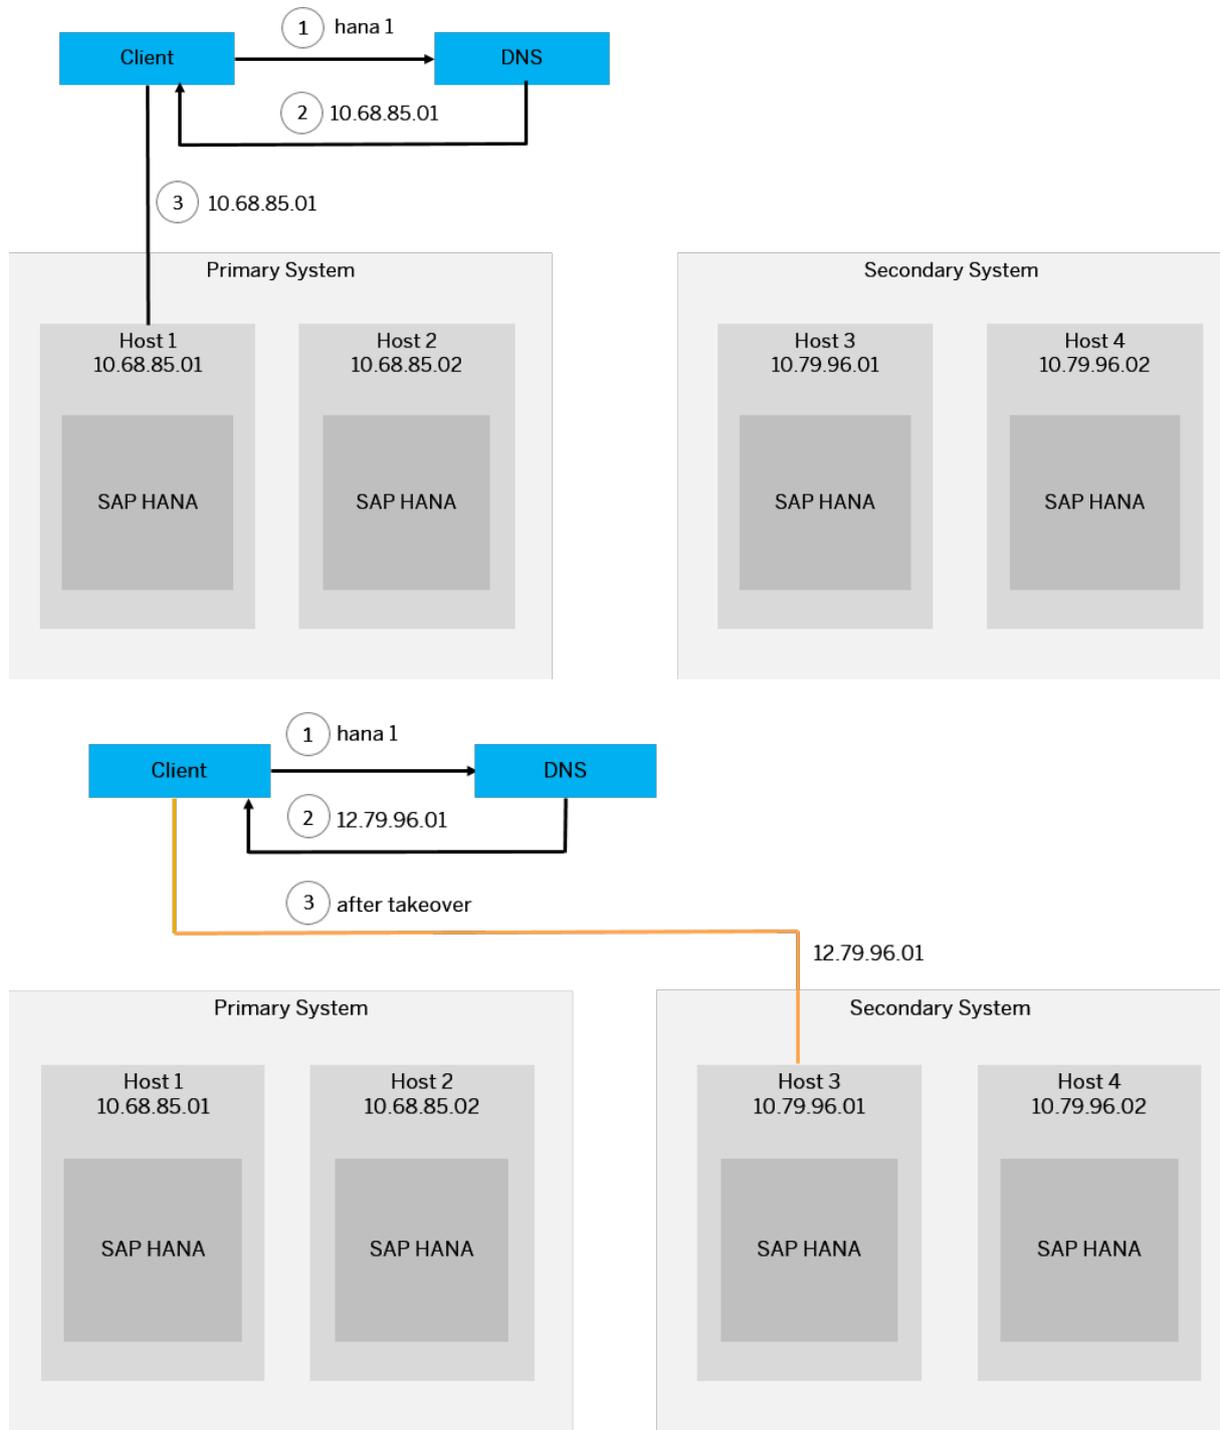
IP redirection can be implemented using a number of actual techniques, for instance with the use of Linux commands which affect the network ARP tables (`ip addr add/del...`), by configuring L2 network switches directly, or by using cluster management software. Following the IP redirection configuration, the ARP caches should be flushed, to provide an almost instantaneous recovery experience to clients.

IP redirection requires that both the primary and failover host(s) are on the same L2 network. If the standby system is in a completely separate L3 network, then DNS redirection is the preferred alternative solution.

Network-based DNS Redirection

DNS redirection is an alternative to IP redirection. DNS is a binding from a logical domain name to an IP address. Clients contact a DNS server to obtain the IP address of the SAP HANA host (step 1 below) they wish to reach. As part of the takeover procedure, a script is executed that changes the DNS name-to-IP mapping from the primary host to the corresponding host in the secondary system (pair-wise for all hosts in the

system). From that point in time, clients are redirected to the failover hosts, as in step 2 of the following diagram:



This solution shares the advantage with IP redirection that there are no client-specific configurations. Further, it supports disaster recovery configurations where the primary and secondary standby systems may be in two completely different network domains (separated by routers). One drawback of this solution is that modifying DNS mappings requires a vendor-proprietary solution. Further, due to DNS caching in nodes (both clients and

intermediate network equipment), it may take a while (up to hours) until the DNS changes are propagated, causing clients to experience downtime despite the recovery of the system.

Cluster Management Software

External cluster management software can be used to perform the client reconnect after takeover. Some of SAP's hardware partners offer an integration of SAP HANA high availability in their cluster management solutions.

Takeover Hooks

They are provided by SAP HANA in the form of a Python script template. Pre- and post-takeover actions are implemented in this script, which are then executed by the name server before or after the takeover. One of the actions could be rebinding the IP addresses.

Related Information

[Implementing a HA/DR Provider \[page 1215\]](#)

[SAP Note 2053504](#) 

10.1.2.8 Performing a Failback

After a takeover has been carried out the roles between primary and secondary can be switched over.

In the case of a failover, the former primary has to be registered as the secondary with the now active primary system. The roles are switched compared to the original setup.

You can perform a failover using the following tools:

- SAP HANA cockpit
For more information, see *Perform a Failback with the SAP HANA Cockpit*.
- SAP HANA studio
For more information, see *Perform a Failback with the SAP HANA Studio*.
- hdbnsutil
For more information, see *Perform a Failback with hdbnsutil*.

Related Information

[Perform a Failback with the SAP HANA Cockpit \[page 1137\]](#)

[Perform a Failback with the SAP HANA Studio \[page 1138\]](#)

[Perform a Failback with hdbnsutil \[page 1140\]](#)

10.1.2.8.1 Perform a Failback with the SAP HANA Cockpit

To perform a failback in the SAP HANA cockpit, register the former primary system as secondary to the current primary.

Context

Use the System Replication tile on the system [Overview](#) page of the former primary to register this system as a new secondary.

Procedure

1. Register the secondary system as follows:
 - a. On the [Overview](#) page of the primary system, choose the [System Replication](#) tile.

The [System Replication](#) page opens. If you performed a data backup before enabling system replication, this page displays overview information on the primary system on the top left and the [Configure System Replication](#) link on the top right.
 - b. Choose [Register Secondary System](#).

The [System Replication Configuration](#) page opens.
 - c. On the [System Replication Configuration](#) page enter the logical name used to represent the secondary system.
 - d. On the [System Replication Configuration](#) page select a replication mode. For more information on the available replication modes, see [Replication Modes for SAP HANA System Replication](#).
 - e. Select an operation modes. For more information on the available operation modes, see [Operation Modes for SAP HANA System Replication](#).
 - f. Enter the host of the source system.

i Note

If you are operating a distributed system on multiple hosts, enter the name of the host on which the master name server is running.

- g. Check [Start Secondary after Registration](#).
2. Review the configured information and choose [Configure](#) on the bottom right.

The [System Replication Configuration](#) dialog opens. After the configuration is complete, the [System Replication Overview](#) page displays information on the secondary site.

Results

The original primary system is now registered as the secondary system with the current primary system (that is, the original secondary system). The secondary system is getting in sync again with the primary system. As such, it is attempting to avoid a full data shipping.

Verify that the secondary system replication status is `All services are active and in sync`.

Related Information

[Replication Modes for SAP HANA System Replication \[page 1093\]](#)

[Operation Modes for SAP HANA System Replication \[page 1094\]](#)

10.1.2.8.2 Perform a Failback with the SAP HANA Studio

You can perform a failback using the SAP HANA studio.

Prerequisites

- You are logged on to both systems as the operating system user (user <sid>adm) or are able to enter these credentials when prompted.
- You have performed a data backup or storage snapshot on the current primary system. In multiple-container systems, the system database and all tenant databases must be backed up.
- The original primary system is not running.
- The current primary system is running.

Context

To fail back to your original primary system, you must switch the roles of your systems back to their original configuration. To do so, the original primary system will have to be started as a secondary system. After both systems are back in sync, you can perform a takeover on the original primary system.

Procedure

1. Register the original primary system as the secondary system as follows:
 - a. In the *Systems* view, right-click the primary system and choose  *Configuration and Monitoring* 
Configure System Replication .

The *Configure System Replication* dialog opens.

i Note

You can also access the *Configure System Replication* dialog from the ► *Landscape* ► *System Replication* ► tab.

- b. Choose *Register Secondary System* and then *Next*.
- c. Enter the required system information and the logical name used to represent the system.

i Note

If you are operating a distributed system on multiple hosts, you enter the name of the host on which the master name server is running.

- d. Specify the log replication mode. For more information on the available replication modes, see *Replication Modes for SAP HANA System Replication*.
- e. Specify the operation mode. For more information on the available operation modes, see *Operation Modes for SAP HANA System Replication*.
- f. Review the configured information and choose *Finish*.
- g. If necessary, start the original primary system.

i Note

The original primary system is started automatically unless you deselected the corresponding option during configuration.

The original primary system is now registered as the secondary system with the current primary system (that is, the original secondary system). As the data that is already available in the original primary system cannot be reused, a complete initialization is carried out. This means that a full data replication takes place until the original primary system is fully in sync.

2. Verify that the secondary system replication status is `All services are active and in sync`. You can see this status in the Administration editor on the *Overview* tab.
3. Fail back to the original primary system as follows:
 - a. In the *Systems* view, right-click the current primary system and choose *Stop System*.
 - b. In the *Systems* view, right-click the original primary system and choose ► *Configuration and Monitoring* ► *Configure System Replication* ►.
 - c. Choose *Perform Takeover* and *Next*.
 - d. Enter the required system information and choose *Next*.
 - e. Review the information and choose *Finish*.
 - f. If necessary, stop the current primary system.

i Note

If the current primary system is still running at the time of takeover, it is stopped automatically unless you deselected the corresponding option during takeover (step 3).

4. Re-register the original secondary as follows:
 - a. In the *Systems* view, right-click the system and choose ► *Configuration and Monitoring* ► *Configure System Replication* ►.

The *Configure System Replication* dialog opens. The *Enable System Replication* option is selected by default.

- b. Choose *Register Secondary System* and then *Next*.
- c. Enter the required system information and the logical name used to represent the secondary system and choose *Next*.
- d. Specify the log replication mode.
- e. Review the configured information and choose *Finish*.
- f. If necessary, start the original secondary system, which is now back in its original role.

i Note

The original secondary system is started automatically unless you deselected the corresponding option during configuration.

Results

The primary system and secondary system have their original roles again.

Related Information

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 1317\]](#)

[Replication Modes for SAP HANA System Replication \[page 1093\]](#)

[Operation Modes for SAP HANA System Replication \[page 1094\]](#)

10.1.2.8.3 Perform a Failback with hdbnsutil

You can perform a failback using `hdbnsutil`.

Context

This is the same procedure as is used for setting up a normal secondary described in *Set Up SAP HANA System Replication with hdbnsutil*. However, in this scenario when the new secondary is registered with the new primary it checks if a delta shipping is possible to re-sync the two sites rather than carrying out a full data shipping. If this is possible it only ships the delta, which significantly reduces the initialization time during registration of the new secondary.

When the new secondary starts up, it checks first if there is a local snapshot available from the time when the system was the primary system. If a snapshot is available the system then checks if it is compatible with the new primary. When both checks are positive the new secondary can be initialized with a delta replica from the new primary.

Related Information

[Set Up SAP HANA System Replication with hdbnsutil \[page 1104\]](#)

10.1.2.9 Disabling SAP HANA System Replication

You can disable SAP HANA system replication for an SAP HANA system by first unregistering all secondary systems and then disabling system replication on the primary system.

You can disable system replication using the following tools:

- SAP HANA cockpit
For more information, see *Disable SAP HANA System Replication with the SAP HANA Cockpit*.
- SAP HANA studio
For more information, see *Disable SAP HANA System Replication with the SAP HANA Studio*.
- hdbnsutil
For more information, see *Disable SAP HANA System Replication with hdbnsutil*.

i Note

There are only three scenarios in which it is necessary to unregister system replication:

- When the secondary system is available, but should be de-coupled permanently
You will be able to use the secondary system as a standard SAP HANA installation afterwards.
- When the secondary system is not available anymore and the primary system needs to be cleaned up in order to be able to register a new system
This can occur when the secondary system was uninstalled or when it cannot be recovered after a disaster.
- When you want to re-establish the original setup after a takeover in a multitier system replication configuration
For more information, see *Restore the Original SAP HANA Multitier System Replication Configuration*.

For more information on the use cases of the command `hdbnsutil -sr_unregister`, see SAP Note 1945676.

Related Information

[Disable SAP HANA System Replication with the SAP HANA Cockpit \[page 1142\]](#)

[Disable SAP HANA System Replication with the SAP HANA Studio \[page 1142\]](#)

[Disable SAP HANA System Replication with hdbnsutil \[page 1143\]](#)

[Restore the Original SAP HANA Multitier System Replication Configuration \[page 1180\]](#)

[SAP Note 1945676](#)

10.1.2.9.1 Disable SAP HANA System Replication with the SAP HANA Cockpit

You can disable SAP HANA system replication in an SAP HANA system using the SAP HANA cockpit.

Prerequisites

The secondary system must be offline.

Context

The *System Replication* tile provides information on the replication mode, the operation mode, the configuration type, and the status of system replication.

Procedure

1. Unregister the secondary system as follows:
 - a. On the *Overview* page of the secondary system, choose the *System Replication* tile.
The *System Replication* tile opens displaying the *System Replication Overview*.
 - b. Choose *Unregister System Replication* on the top right.
2. Disable system replication on the primary system as follows:
 - a. On the *Overview* page of the primary system, choose the *System Replication* tile.
The *System Replication* tile opens displaying the *System Replication Overview*.
 - b. Choose *Disable System Replication* on the top right.

10.1.2.9.2 Disable SAP HANA System Replication with the SAP HANA Studio

You can disable SAP HANA system replication in an SAP HANA system using the SAP HANA studio.

Prerequisites

- You are logged on to both systems as the operating system user (user <sid>adm) or are able to enter these credentials when prompted.

- The secondary system must be offline.

Procedure

1. Unregister the secondary system as follows:
 - a. In the *Systems* view, right-click the primary system and choose ► *Configuration and Monitoring* ► *Configure System Replication* ▾.
The *Configure System Replication* dialog opens.

i Note

You can also access the *Configure System Replication* dialog from the ► *Landscape* ► *System Replication* ▾ tab.

- b. Choose *Unregister secondary system* and then *Next*.
 - c. Enter the required system information and choose *Next*.
 - d. Review the configured information and choose *Finish*.
2. Disable system replication on the primary system as follows:
 - a. In the *Systems* view, right-click the primary system and choose ► *Configuration and Monitoring* ► *Configure System Replication* ▾.
 - b. Choose *Disable system replication* and choose *Next*.
 - c. Review the configured information and choose *Finish*.

10.1.2.9.3 Disable SAP HANA System Replication with hdbnsutil

You can disable SAP HANA system replication in an SAP HANA system with hdbnsutil.

Prerequisites

- You are logged on to both systems as the operating system user (user <sid>adm) or are able to enter these credentials when prompted.
- The secondary system must be offline.

Procedure

1. Stop the secondary system and unregister it as follows:

```
hdbnsutil -sr_unregister
```

For other use cases of the command `hdbnsutil -sr_unregister`, see SAP Note 1945676.

If system replication is out of sync and you need to register again the initial secondary system, use the command `hdbnsutil -sr_register`. It is not needed to unregister the secondary system before registering it again.

2. Disable system replication on the primary system as follows:

```
hdbnsutil -sr_disable
```

Related Information

[SAP Note 1945676](#)

10.1.2.10 Resync Optimization

Whenever the primary and the secondary systems are disconnected, SAP HANA system replication is out of sync. To get in sync again, a shipping of the missing data is initiated.

The system tries to avoid a full data shipping and to achieve a resync with a delta data or a log shipping. To get the primary and the secondary systems in sync again, their persistencies (that is, the data and log volumes) must be compatible. The system that is to be registered as the secondary system checks if its persistence is compatible with the primary system. If this check succeeds, a delta shipping can be carried out instead of requesting a full data shipping from the primary system. There is a maximum of three checks executed by the secondary system in the following order:

1. Check if the newest savepoint is compatible:
 - The to-be secondary system checks if its newest savepoint is compatible.
 - This check most likely succeeds if the secondary system has just been shut down for a short time.
2. Check if the newest replication snapshot is compatible:
 - Replication snapshots are written on the system replication primary and secondary sites while replication is up and running
 - They are created on the secondary site each time a savepoint is written.
 - They are created periodically on the primary site (time and volume based) to preserve a state that is known to be shipped to the secondary site. As the snapshot verification takes some time, a replication snapshot that is not yet verified to be shipped may have been created on the primary system.
 - This check most likely succeeds after a test takeover on the secondary system because this state has to be available also on the primary system.
3. Check if the active replication snapshot is compatible:
 - The active replication snapshot is a special replication snapshot, created on a primary site and verified to be shipped to the secondary site.
 - This check most likely succeeds during a failback operation because it is created on the old primary and the snapshot is verified to be shipped.

The first savepoint or snapshot that is compatible with the primary site will be used for delta data shipping. If none of the three savepoints or snapshots are compatible, then a full data shipping will automatically be carried out.

i Note

If system replication is out of sync and you need to register again the initial secondary system, use the command `hdbnsutil -sr_register`. It is not needed to unregister the secondary system before registering it again. Unregistering the initial secondary system before registering it would hinder an optimized resync and would trigger a full data shipping.

Depending on the chosen operation mode, two different techniques are in place to achieve a resync: data retention and log retention. For more information, see *Data Retention* or *Log Retention*

Related Information

[Data Retention \[page 1145\]](#)

[Log Retention \[page 1145\]](#)

10.1.2.10.1 Data Retention

Whenever the primary and the secondary systems are disconnected SAP HANA system replication is out of sync. To get in sync again, a delta shipping of the missing data is initiated.

In the `delta_datashipping` operation mode, the primary system sends the incremental data to resync after a disconnect if the last snapshot that was successfully sent to the secondary system is still available. If it is no longer available, a full set of data is necessary to get in sync again.

10.1.2.10.2 Log Retention

With the operation modes `logreplay` and `logreplay_readaccess`, log segments can be marked as `retained` so that they can sync a secondary system after a disconnect.

With continuous log replay, delta data shipping cannot be used to sync a secondary site anymore. This is because although the primary and secondary persistence is logically compatible they are no longer physically compatible. This means the data, that is contained in the persistence is the same, but the layout of the data on pages can be different on the secondary site. Therefore a secondary site can sync via delta log shipping only. This is relevant for the following use cases:

- The secondary site has been disconnected for some time (for example, because of a network problem or temporary shutdown of secondary site)
- A former primary site has been registered for failback

The secondary site only uses log of the online log area of the primary SAP HANA system for syncing. The log must be retained for a longer time period than before to be able to sync the secondary site. If syncing via delta

log shipping does not work, for example because the log has been reused, a full data shipping becomes necessary. To avoid this, if possible, the concept of Log Retention has been introduced.

Log Retention for Secondary Disconnect (on primary site)

When a secondary system configured with the operation mode `logreplay` or `logreplay_readaccess` is disconnected, the primary system should not reuse the log segments in the online log area that are required to sync the secondary site via delta log shipping. These log segments are marked as `RetainedFree` until the secondary has successfully synced again. If a secondary system is stopped, the log volume will grow on primary site, until the log volume has filled up with log segments. Once the secondary system reconnects and has synced the missing log, these log segments are then set to `Free` and can be reused after that.

Log segments are retained on the primary as long as the secondary site is registered, but not connected to the primary site. Therefore, if a secondary site is shut down and not used for a longer time period unregister it first, to prevent log volumes from filling up on the primary site. However, in this case a full data shipping will be necessary when the system reconnects. This behavior is automatically turned on, if a secondary system with operation mode `logreplay` or `logreplay_readaccess` is registered.

Log Retention for Failback (on secondary site)

On the secondary site, log retention is required to do a failback with optimized data transport. The primary site periodically creates persistence snapshots during replication. After takeover, when the old primary is started again as secondary, the most recent snapshot is opened on the old primary site and the missing log is requested from the new primary..

With respect to log retention we have to distinguish between two situations:

1. Log Retention During Replication
During replication time the secondary site keeps all log starting from the last primary snapshot position. Old log is automatically released after a new snapshot has been created on the primary site. This behavior is turned on by default and it ensures that during replication only a few `RetainedFree` segments are kept online. They are needed to fill the gap between the primary snapshot and the current potential takeover log position.
2. Log Retention After Takeover
After takeover the new primary has to keep log until a new secondary site is registered and has synced the missing log. Because syncing can take some time this behavior has to be explicitly turned on by setting `global.ini/[system_replication]/enable_log_retention = on`

After the new secondary has been connected, the behavior will be the same as described in the previous section.

If you have a setup in which there will be frequent failbacks between two sites, we recommend that you set the following parameter on both sites to simplify configuration: `global.ini/[system_replication]/enable_log_retention = on`

In this case, no additional configuration change is required, when sites are being switched.

Log Retention and Disk Full

The parameter can be used to specify how the SAP HANA system behaves when many log segments of the type `RetainedFree` are created.

If `logshipping_max_retention_size` has been set to a value other than 0, when no secondary is connected log segments are not reused. This occurs even if they are truncated and backed up until the max size limit has been reached or the system runs into a log full situation. If the max size limit is reached or in a log full situation, segments that are only kept for syncing the secondary site are reused. This setting prevents the system from a standstill on the primary site due to too many log segments, which are held for syncing the secondary site. With this setting the primary is kept running with the drawback that the secondary cannot sync anymore via delta log shipping. In this case a full data shipping will become necessary (soft limit).

If `logshipping_max_retention_size` is configured to 0, then log segments required for secondary syncing are not reused and a log full results in a system standstill on the primary site until log writing can continue. With this setting, being able to sync the secondary has priority over standstill on the primary. When the reason for the log full has been resolved (on the primary or secondary site), transaction processing can continue (hard limit).

i Note

The default setting `logshipping_max_retention_size = 1048576` (MB) of 1 TB means that 1 TB of size is configured for every service, which replicates data to a secondary system (that is, every service owning a persistence in form of data and log volume).

🔗 Example

If the services `nameserver`, two `indexservers` (for example, two tenant databases) and an `xsengine` are running in your SAP HANA system, the total configured log retention size will be 4 TB (4 x 1 TB). With this setting it can happen that the disk full is reached before the `RetainedFree` marked log segments are overwritten.

If you want to change the default value of 1 TB, you can do this in the `global.ini`. Another option is to set this parameter in the service ini files individually. For example, if the value is set in the `global.ini` of the `SystemDB`, in the `global.ini` of a tenant database, and in the `indexserver.ini` of a tenant database, the `indexserver.ini` setting would win and will be taken for log retention of this `indexserver`.

Log Retention and Multitarget System Replication

When the primary system replicates data changes to more than one secondary system, you should use force log retention and log retention propagation to reach an optimized re-sync and avoid a full data shipping after takeover or other disconnect situations..

Force log retention is used on a system to retain log until it is actively disabled. To use force log retention, enter the value `force_on_takeover` for the `enable_log_retention` configuration parameter.

If `enable_log_retention = force_on_takeover` is configured, the log will be retained during replication for all direct secondaries until a takeover is executed. During takeover, the parameter is set to `force`. This means the log will be retained independently of any secondary system.

❁ Example

A typical usage scenario is described in the following steps:

1. Configure all systems with `[system_replication]/enable_log_retention = force_on_takeover`
2. During takeover on a secondary system, if `force_on_takeover` is set, the value is changed to `enable_log_retention = force`. This means that starting from the takeover, the log is retained until it is explicitly disabled.
3. Re-register all required systems until the landscape is again fully functional.
4. Reset `[system_replication]/enable_log_retention = force_on_takeover` on the system on which takeover has been executed before re-establishing the original configuration.

The configuration must be done manually (for example, by the administrator or using setup scripts) because the SAP HANA system doesn't know when the system landscape has been completely reconfigured.

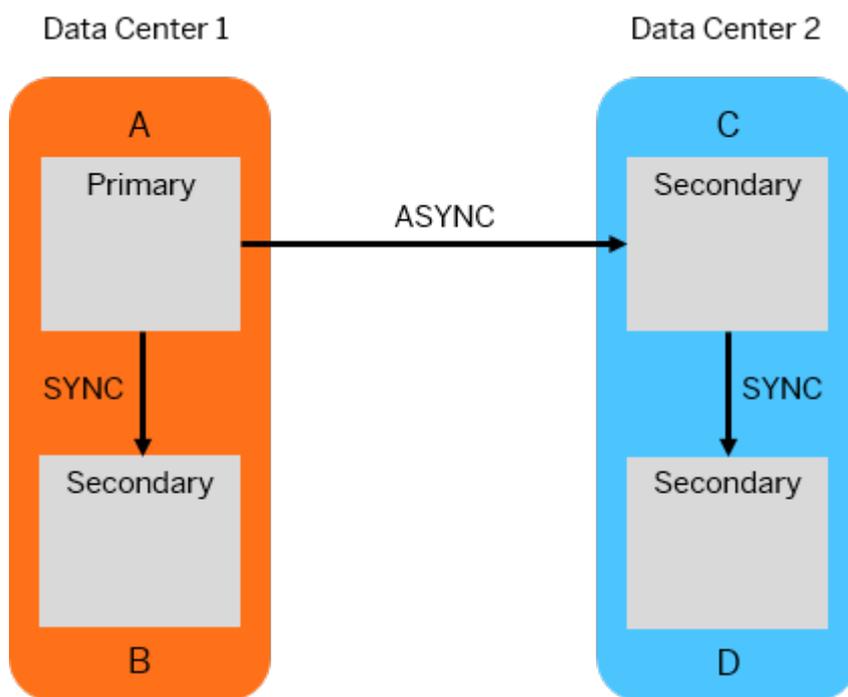
Log retention propagation is used to retain log based on the smallest savepoint log position in the whole system replication landscape. Log retention propagation should be enabled if you want to re-order your systems in a complex setup. This can be done by setting the following parameter in `global.ini`: `[system_replication]/propagate_log_retention = on`. If you want to propagate log retention in a system replication landscape between all systems, this parameter should be set on all systems in the landscape.

When you set this parameter on a system, it behaves as follows:

- It sends the minimum log position of its own savepoint and the retain log position it gets from all direct secondaries as retain position to its source system
- It sends the minimum log position of its own savepoint and the retain log position it gets from all direct secondaries to the secondaries as retain log position
- It uses the minimum log position it gets from all direct secondaries and its source system (if not primary) as own retain log position

To explain these concepts we are using the setup described in *Multitarget System Replication*. In this setup, primary system A replicates data changes to secondary system B located in the same data center. Primary system A also replicates data changes to the secondary system C located in another data center. Secondary

system C is a source system for a further secondary system D located in the same data center with system C. For a quick overview, use the graphic below:



❖ Example

If there is a takeover on secondary system B, you must register system C to B and system A to B to re-setup the original configuration. To avoid a full data shipping for both systems, system B must retain all the log until systems A and C have synced again. This can't be accomplished by setting `global.ini/[system_replication]/enable_log_retention = on` because system B doesn't know how many systems must be re-attached until the landscape is back in its functional state.

Force log retention should be used on system B until systems A and C are registered again and synced.

If you want to re-order your systems, enable log retention propagation. Log retention without propagation only respects the direct neighbors. For example, if system D is stopped in this setup, system C retains log for D, but not for A and B. If system D is re-attached to systems B or A and propagation is not turned on, log could be missing because systems A and B do not retain their log with respect to D.

Related Information

[Multitarget System Replication \[page 1153\]](#)

[Disaster Recovery Scenarios for Multitarget System Replication \[page 1154\]](#)

10.1.2.11 Secondary Time Travel

You can start the secondary system in online mode on a previous point in time.

Secondary time travel is a special version of a takeover and can be used to quickly access again data, which was deleted in the original system.

i Note

You can use secondary time travel only with the operation modes `logreplay` or `logreplay_readaccess`. Configuring the SYNC replication mode with the full sync option, does not lead to a freeze of the primary system upon takeover.

To prepare the secondary system for time travel, snapshots are kept on the secondary system for a defined time travel period. These snapshots can be used later to start the system on an older point in time. Additional log will be retained on the secondary system starting from the oldest time travel snapshot. After opening the old snapshot, the additional log has to be replayed to reach the requested point in time.

You can start the secondary system on an older point in time using `hdbnsutil -sr_timetravel`. During the execution of `hdbnsutil -sr_timetravel`, the specified time and location are stored internally in a dedicated file in the SAP HANA directory. When calling `hdbnsutil -sr_timetravel`, it can be specified if takeover hooks should be called. If the parameter is not explicitly specified, the default value from the configuration parameter `timetravel_call_takeover_hooks` will be used.

The secondary system will enter in online mode on the specified point in time during restart. After restart, the other services read the requested point in time and open their persistence using this information. If the requested point in time cannot be reached, then time travel will be aborted. A check ensures that there are time travel snapshots older than the start time for each service.

You can monitor the start time or log position of the system using `M_SYSTEM_REPLICATION_TAKEOVER_HISTORY`.

For more information on how to execute secondary time travel, see *Execute Secondary Time Travel*.

Related Information

[Execute Secondary Time Travel \[page 1151\]](#)

[Configuration Parameters \[page 1152\]](#)

10.1.2.11.1 Execute Secondary Time Travel

You can start the secondary system in online mode at a previous point in time.

Prerequisites

- Set the `global.ini/[system_replication]/timetravel_max_retention_time` parameter to define the time period to which the secondary system can be brought back in the past.
- You can set the `global.ini/[system_replication]/timetravel_snapshot_creation_interval` parameter to adjust the secondary snapshot creation.

Note

Set the parameters carefully to avoid log full or disk full situations. For time travel to work, log and snapshots are kept online in the data area. Because of this, log and data will grow on the secondary system when time travel is turned on. The system workload determines how much data is needed. For a full list of available parameters, see *Configuration Parameters*.

Context

After setting the parameters, the secondary system starts retaining log and keeping created snapshots. After retaining sufficient log and data, the secondary system is ready for time travel.

You can start the secondary system on an older point in time as follows:

Procedure

1. Stop the SAP HANA database.
2. Execute `hdbnsutil -sr_timetravel --startTime=<startTime> [--callTakeoverHooks=on|off]`.

For `startTime` use the following format specified in UTC: `dd.mm.yyyy - hh.mm.ss`

3. Start the SAP HANA database.

Related Information

[Configuration Parameters \[page 1152\]](#)

10.1.2.11.2 Configuration Parameters

Several parameters are available for configuring secondary time travel.

Use the following parameters to configure secondary time travelling. The parameters are defined in the `system_replication` section . All parameters are set on the secondary system.

Parameter	timetravel_max_retention_time
Type	integer
Unit	minutes
Default	0
Description	If set to 0, secondary time travel is turned off. If this parameter is set to a value different from 0, the secondary system can be brought online up to the defined time period in the past.

Parameter	timetravel_snapshot_creation_interval
Type	integer
Unit	minutes
Default	1440
Description	<p>Defines how frequently snapshots are created for secondary time travel. Time travel snapshots are kept until they get older than the defined <code>timetravel_max_retention_time</code> parameter. If a takeover needs to be done on an older point in time, the snapshot that best fits the requested point in time will be opened and the remaining changes are applied via <code>logreplay</code>.</p> <p>A new snapshot is created when the time period defined in this parameter has passed since the last snapshot creation. Snapshots older than the time period defined in <code>time_travel_max_retention_time</code> are dropped.</p>

Parameter	timetravel_call_takeover_hooks
Type	bool
Values	true, false
Default	false
Description	Indicates if takeover hooks should be called during secondary time travel.

10.1.2.12 Multitarget System Replication

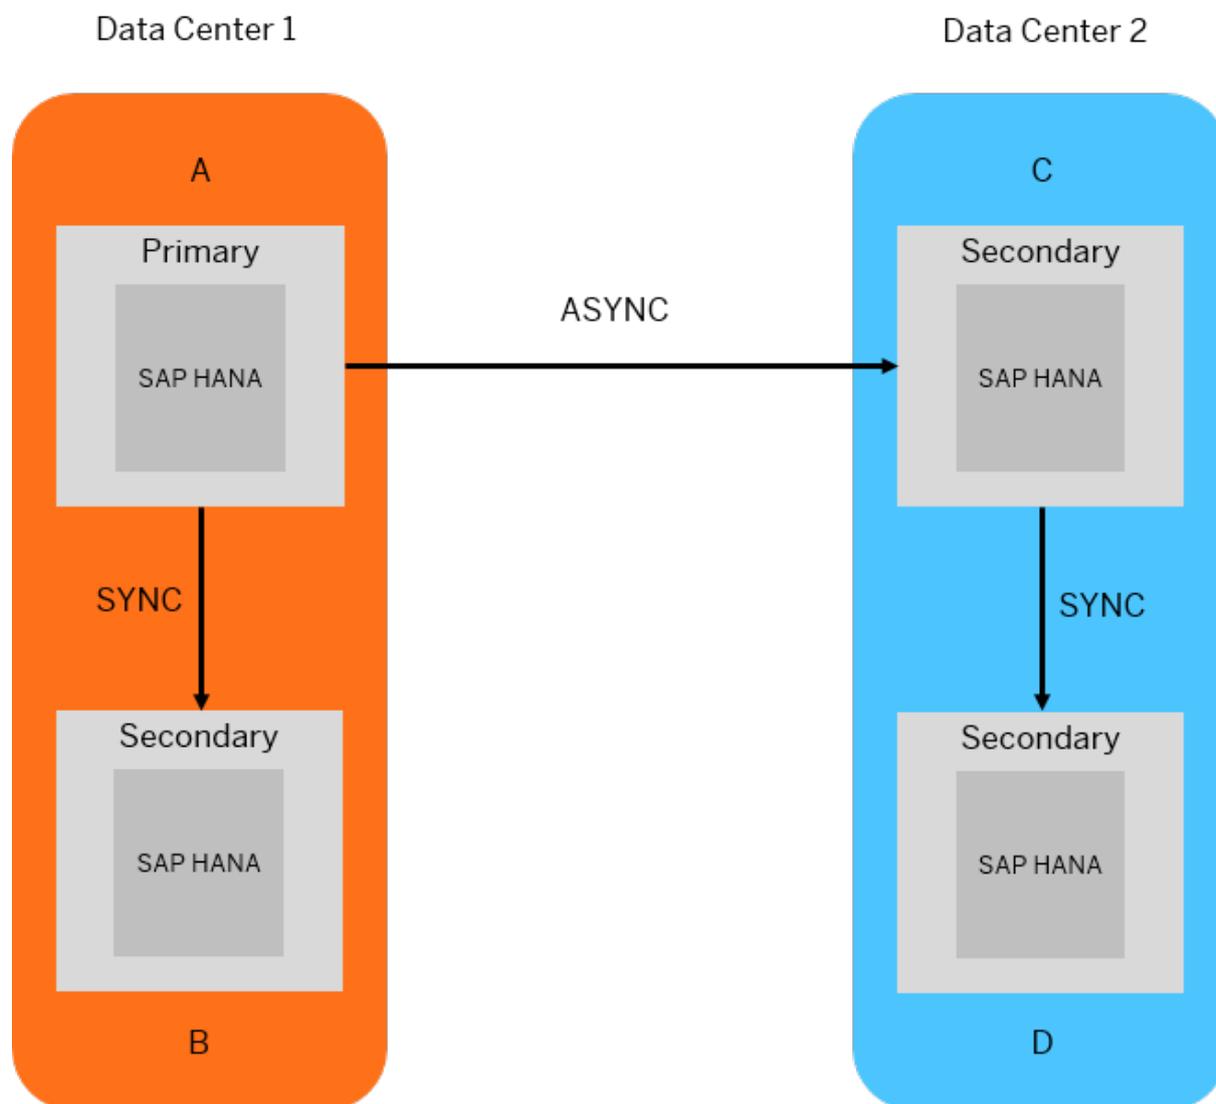
With multitarget system replication the primary system can replicate data changes to more than one secondary system.

Multitarget system replication can bring advantages for several use cases:

- Update scenarios
- Rearrangements of system replication multiter chains
- Reaching higher availability (before stopping existing structures, new structures can be build and established)

The graphic below visualizes a possible setup for multitarget system replication.

Primary system A in data center 1 replicates data changes to secondary system B in the same data center. Primary system A also replicates data changes to secondary system C in data center 2. Secondary system C is a source system for a further secondary system D located in the same data center with system C.



❖ Example

To configure a multitarget system replication with the setup displayed in the graphic, follow these steps:

1. On primary system A in data center 1:
Create backups and enable system replication.
2. On the local secondary system B in data center 1:
Stop the system, register it to A, and start the system.
3. On the remote secondary system C in data center 2:
Stop the system, register it to A, and start the system.
When secondary system C is online, set this system as source system to serve another secondary.
4. On the remote secondary system D in data center 2:
Stop the system, register it to C, and start the system.

To understand how to handle in different disaster recovery scenarios, see *Disaster Recovery Scenarios for Multitarget System Replication*.

Related Information

[Operation Modes for SAP HANA System Replication \[page 1094\]](#)

[Disaster Recovery Scenarios for Multitarget System Replication \[page 1154\]](#)

[Full Sync Option for SAP HANA System Replication \[page 1108\]](#)

[Log Retention \[page 1145\]](#)

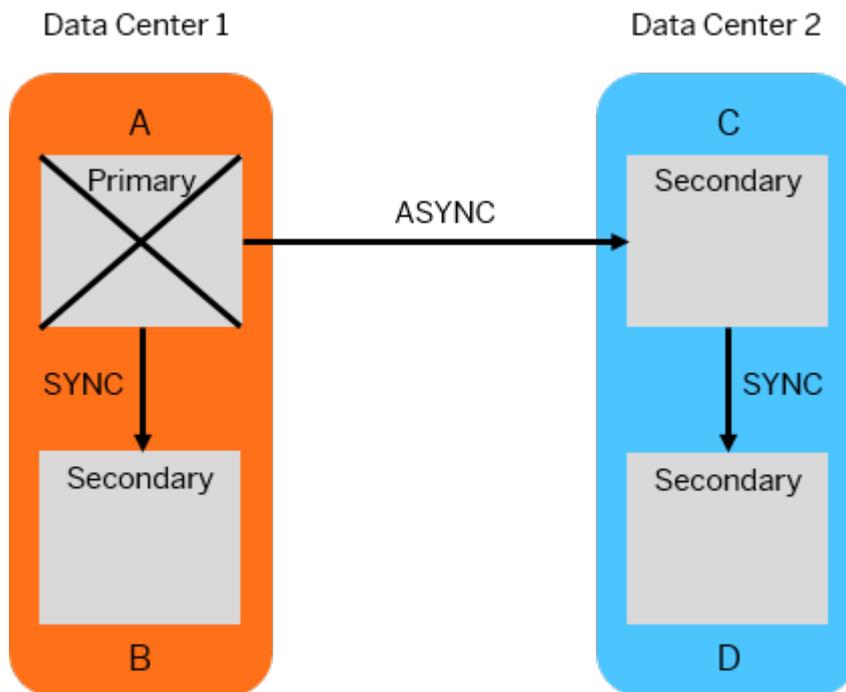
[Configure Secure Communication \(TLS/SSL\) Between Primary and Secondary Sites \[page 1206\]](#)

10.1.2.12.1 Disaster Recovery Scenarios for Multitarget System Replication

Several solutions are available when the systems involved in a multitarget system replication configuration fail.

We are using the setup described in *Multitarget System Replication* to exemplify the procedure. In this setup, primary system A replicates data changes to secondary system B located in the same data center. Primary system A also replicates data changes to the secondary system C located in data center 2. Secondary system C is a source system for a further secondary system D located in the same data center with system C.

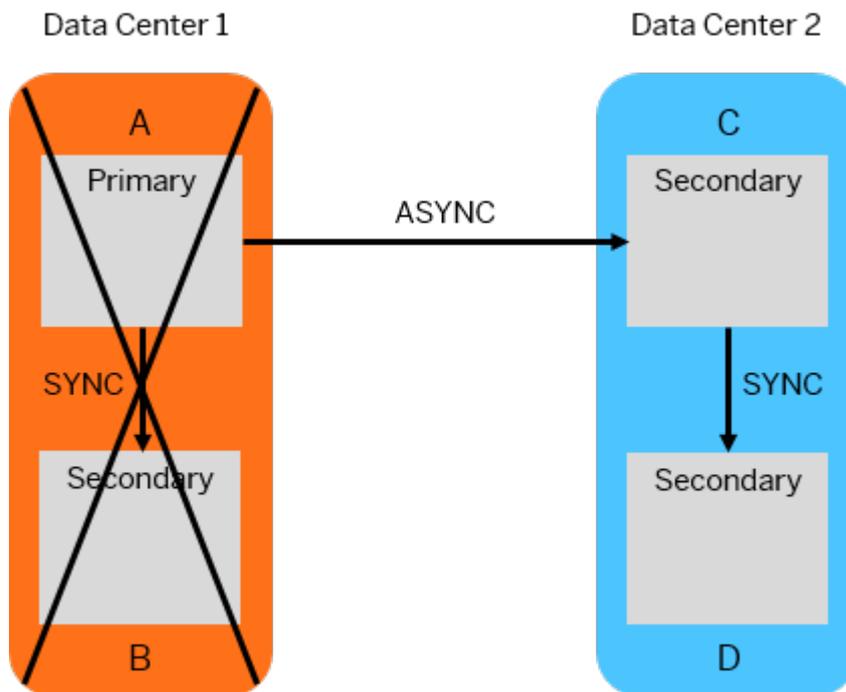
Failure on Primary System A



When primary system A fails, proceed as follows:

1. Take over on secondary system B in data center 1.
2. Register secondary system C in data center 2 to the new primary system B in data center 1. Then, register secondary system D in data center 2 to secondary system C.
3. After the failure on the previous primary system A is solved, register it to the new primary system B in data center 1.

Failure of Data Center 1



When all the systems in data center 1 fail, proceed as follows:

1. Take over on secondary system C in data center 2.
2. After the failure on the previous primary system is solved, register system A to the new primary system C in data center 2.
3. Register secondary system B as tier 3 to system A in data center 1.

For more information about takeover and failback, see *Performing a Takeover* and *Performing a Failback*.

Related Information

[Performing a Takeover \[page 1126\]](#)

[Performing a Failback \[page 1136\]](#)

[Full Sync Option for SAP HANA System Replication \[page 1108\]](#)

[Log Retention \[page 1145\]](#)

10.1.2.13 Active/Active (Read Enabled)

Active/Active (read enabled) enables SAP HANA system replication to support read access on the secondary system.

Active/Active (Read Enabled)

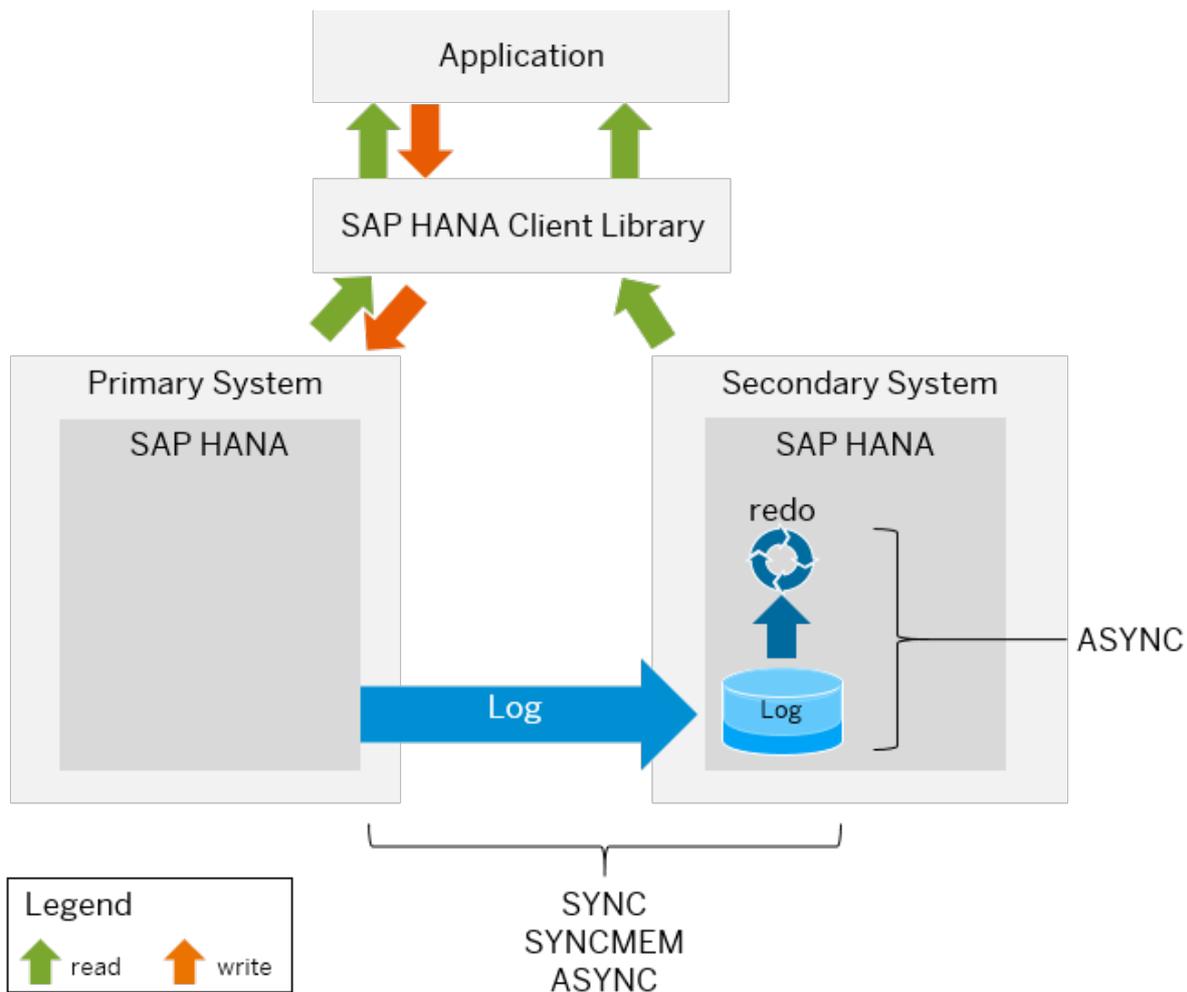
Active/Active (read enabled) reduces the load on the primary system but does not double the capacity; it simply extends read capabilities. In an Active/Active (read enabled) system replication configuration, the SQL ports on the secondary system are open for read access. This makes it possible to use the secondary system for read-intense tasks and to have a better balance of workloads improving the overall performance of the SAP HANA database.

Active/Active (read enabled) is integrated into the System Replication solution and gets activated with the operation mode `logreplay_readaccess`. This operation mode provides fast takeovers, reduced need for bandwidth in continuous operation, and support for replication modes such as SYNC (with or without the full sync option), SYNCMEM, and ASYNC. The redo log replay runs asynchronously to the primary operations.

i Note

In case of an outage, all functions concentrate on the secondary system. Because of this, the sizing of the secondary system is important to ensure the right performance in disaster scenarios.

The following graphic focuses on the Active/Active (read enabled) configuration. The primary system is fully active and supports reading and writing, while the secondary system is enabled for read queries. For a detailed illustration of the general system replication processes, see *System Replication*.



Operation Mode `logreplay_readaccess`

Using this operation mode while configuring your system replication, read access becomes possible on the secondary system by establishing a direct connection to the secondary system or by providing a `SELECT` statement from the primary system with a `HINT`. For more information, see also *Client Support for Active/Active (Read Enabled)* and *SAP HANA SQL and System Views Reference*.

To use this operation mode, the primary and the secondary systems must have the same SAP HANA version. For this reason, read-only access to the secondary system is not possible during a rolling upgrade until both versions are the same again.

An active secondary system is only supported for tier 2. In a 3-tier system using Active/Active (read enabled), the `logreplay_readaccess` mode is required between the primary and the active secondary systems, while the `logreplay` mode is required between the other (tier 2 and tier 3) secondary systems.

i Note

With the operation mode `logreplay_readaccess` the secondary allows read-only access on column tables via SQL providing a delayed view on the data compared to the primary system. There is no minimum delay guarantee. The read access on system and monitoring views is supported as well.

For more information about the operation mode `logreplay_readaccess`, see also *Operation Modes for SAP HANA System Replication*.

License Management

In an Active/Active (read enabled) configuration, the secondary system is operated automatically with the license key of the primary system. Changes of the license key are done on the primary system and replicated to the secondary system. For more information, see *Managing SAP HANA Licenses*.

Related Information

[System Replication \[page 1085\]](#)

[Setting Up SAP HANA System Replication \[page 1098\]](#)

[Operation Modes for SAP HANA System Replication \[page 1094\]](#)

[SAP HANA SQL and System Views Reference](#)

[Managing SAP HANA Licenses \[page 305\]](#)

<https://www.youtube.com/watch?v=dWQYGOi4c7g> 

[SAP Note 2391079](#) 

10.1.2.13.1 Generic Conditions for Active/Active (Read Enabled)

When using the secondary system for read access, several aspects need to be considered.

Points to Consider

- The processors in the primary and secondary systems must be both either Intel-based or IBM Power-based with the same byte ordering. A platform mixture is not supported.
- The secondary system allows read access if the primary system runs the same SAP HANA version. A different version leads to prohibiting the read access until the same software version is used.
- The redo log replay runs as an asynchronous process on the secondary system. The secondary system provides statement level snapshot isolation with potentially delayed view on the data and no minimum delay guarantee.

- The secondary system gets its own virtual IP addresses or host names representing the secondary function.
- The query execution in the secondary system is rejected if it needs background migrations requiring redo log writes (e.g. L2-Delta migration).

→ Recommendation

Perform the migration in the primary system (e.g. load table) and wait until it is replicated and replayed in the secondary system.

- Internal processes for operations like ColumnStore delta merges take place on the secondary system.
- DML execution on an Active/Active (read enabled) secondary system is possible for row store no-logging retention tables and global temporary tables. Their Explain Plan is also available on the secondary system. For more information, see *Data Manipulation Statements*.

Limitations

- Active/Active (read enabled) is supported in a multitier SAP HANA replication system. However, read access is limited to tier 2. The `logreplay` operation mode is required between tier 2 and the further tier level and no read access connections can be opened to the further tier level.
- If Active/Active (read enabled) is used with Dynamic Tiering services, there is no read access to Dynamic Tiering data on the secondary system.
- The export of tables is possible with CSV as target. However, binary exports on the secondary system are not supported.
- In a multitarget system replication only one secondary system allows read access.

Support for Multiple SAP HANA Databases

It is possible to use the read-enabled secondary system for other SAP HANA systems such as development or QA environments. In this case the following sizing conditions apply:

- The secondary hardware must offer the same CPU and memory capacities as those offered by the primary system **plus** the resources for the additional system.
- After a takeover, the system must be capable of handling both the primary's writing load and the secondary's reporting load.

For more information about this scenario, see also SAP Note 1681092 *Multiple SAP HANA DBMSs (SIDs) on one SAP HANA system*.

Related Information

[SAP Note 1681092](#)

[SAP Note 2447994](#)

10.1.2.13.2 Connection Types

Connecting to an Active/Active (read enabled) system allows you to take advantage of a secondary system for better overall performance.

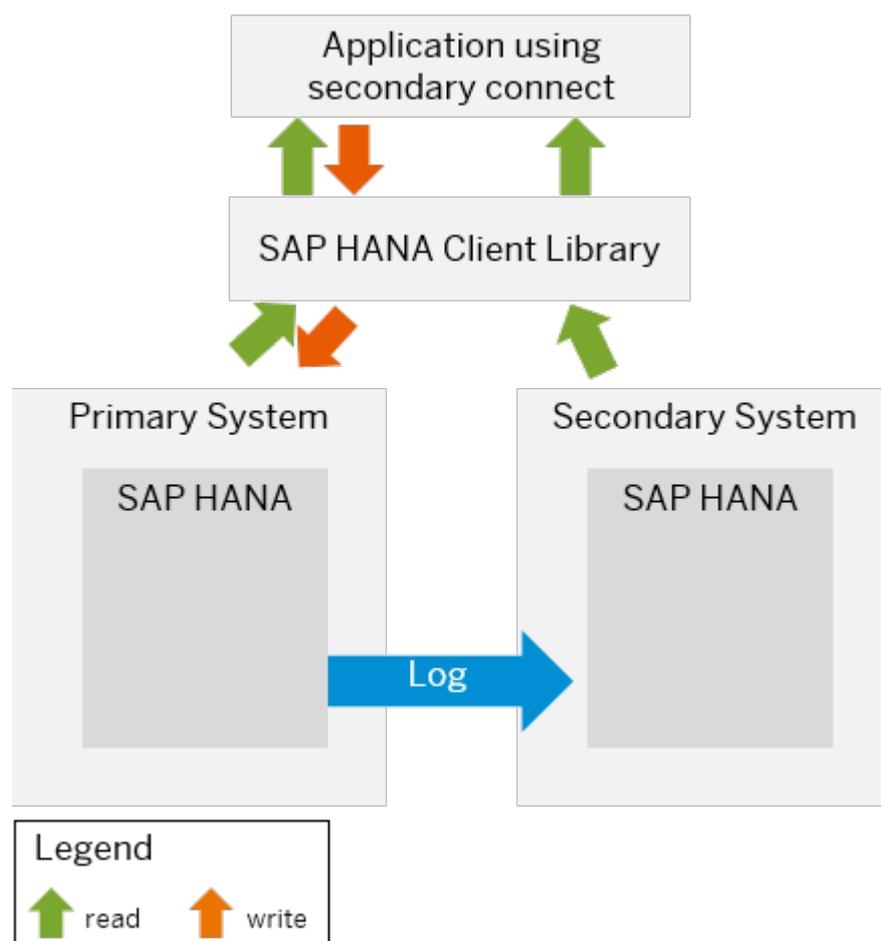
There are two ways to execute SQL queries on the read enabled secondary system:

- Opening an explicit connection to the secondary system
- Executing a SQL statement on the primary system which is redirected to the secondary system according to a hint

Explicit read-only connection to the secondary system

For this connection type the application opens the connection to the secondary system. There is no session property sharing.

The following graphic illustrates an explicit read-only connection to the secondary system:



Implicit hint-based statement routing

You can pass the SQL query to the primary system and add a hint saying that this statement should be preferably executed on the read enabled secondary system.

i Note

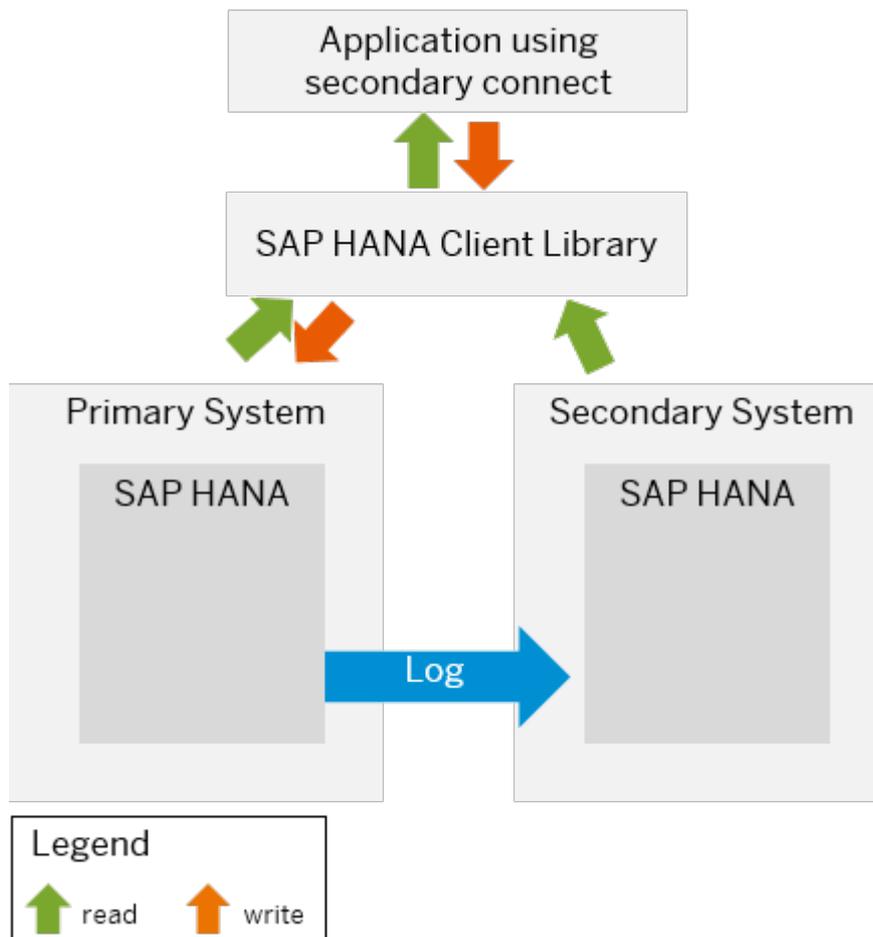
Hint-based statement routing is supported by SAP HANA ODBC, SQLDBC, ADO.Net, JDBC drivers for SAP HANA 2.0 and the SAP HANA Node.js.

The SAP HANA client opens an additional connection to the secondary system according to the host information returned by the primary system.

This connection type unfolds as follows:

1. The SAP HANA client sends the statement-prepare with hint to the primary system.
2. The primary system decides where to execute the statement and returns the result to the SAP HANA client.
3. The SAP HANA client sends the statement execution call to the secondary system. The session property changes are handed over to the secondary system via the SAP HANA client. If the secondary system cannot execute the statement, it returns an error and the SAP HANA client sends the statement to the primary system.

The following graphic illustrates an implicit hint-based statement routing:



For more information on the hint-based statement routing, see *Hint-Based Routing for Active/Active (Read Enabled)*.

To cancel long-running sessions on a read enabled secondary system, use: `ALTER SYSTEM CANCEL SESSION.`

Limitations:

- The SAP HANA client has an open working connection to the primary system.
- Hint-based statement routing is not supported for CALL procedures.
- Temporary tables are not supported.
- Hint-based statement routing is not supported for write transactions.

10.1.2.13.3 Configuration Parameters

Several parameters are available for configuring Active/Active (read enabled).

Parameter: `operation_mode`

Values: `logreplay_readaccess`

System: `Secondary`

Description: System Replication uses an initial data shipping to initialize the secondary system. After that, only log shipping is done and the log buffers received by the secondary system are being replayed there. Savepoints are executed individually for each service. Column table merges are executed on the secondary system. Additionally, read access via SQL is provided to the secondary system.

Relevant for Active/Active (read enabled) are also `enable_log_retention` and `logshipping_max_retention_size`. For more information about these parameters, see *SAP HANA System Replication Configuration Parameters*.

Related Information

[SAP HANA System Replication Configuration Parameters \[page 1109\]](#)

10.1.2.13.4 Checking the Active/Active (Read Enabled) Configuration

You can check if your system replication is configured as an Active/Active (read enabled) system.

Use the following tools to check if your system replication is configured as an Active/Active (read enabled) system:

SAP HANA Cockpit

On the *System Replication* tile on the primary system, the `logreplay_readaccess` operation mode indicates that your system is an Active/Active (read enabled) system. Additionally, an enabled secondary read access informs you that the SQL ports are open for reading on the Active/Active (read enabled) secondary system.

On the system overview page of the secondary system, the *Mode: read-only* indicates that your system is an Active/Active (read enabled) system. Additionally, the *Delay* in ms is shown on top indicating how far behind the consistent view on the data of this secondary system is compared to the current data of the primary system.

SAP HANA Studio

On the primary system, select `M_SYSTEM_REPLICATION` from the monitoring view. To find out if your system is an Active/Active (read enabled) system, verify the columns *OPERATION_MODE* and *SECONDARY_READ_ACCESS_STATUS*.

Command Line

As `<sid>adm` run one of the following commands and then look for the operation mode `logreplay_readaccess`:

```
python $DIR_INSTANCE/exe/python_support/systemReplicationStatus.py --
sapcontrol=1 | grep OPERATION_MODE
service/ld4144/30207/OPERATION_MODE=logreplay_readaccess
service/ld4144/30201/OPERATION_MODE=logreplay_readaccess
service/ld4144/30203/OPERATION_MODE=logreplay_readaccess
```

or

```
hdbnsutil -sr_state | grep "operation mode"
operation mode: logreplay_readaccess
```

10.1.2.13.5 Memory Management

Several parameters can be used to set the memory limit for read accesses on the secondary system.

The total statement memory is limited to 50 % of the global allocation limit, because 50% of the storage is reserved for logreplay. Logreplay should not fail because of memory limitations.

Use the parameters below to set the memory limit for read accesses on the secondary system:

Section: memorymanager

Parameter: `sr_total_statement_memory_limit`

Type: int (GB)

Default: (empty)

Description: Memory limit in GB:

- (empty): 50% of global allocation limit
- 0: disable the feature
- N: set the value as a limit

Section: resource_tracking

Parameter: `sr_memory_tracking`

Type: bool

Default: on

Description: Enables/disables memory tracking on the secondary system

Parameter: `sr_memory_tracking`

Type: bool

Default: on

Parameter: `sr_memory_tracking`

Description: Enables/disables memory tracking on the secondary system

10.1.2.13.6 Authentication Methods

There are several authentication methods supported for an Active/Active (read enabled) system replication.

The following authentication methods are supported for the primary system:

- Basic (User Name/Password)
- Kerberos
- SAML
- Session Cookies

The secondary system delegates the authentication phase to the primary system using the existing communication channel from the secondary system to the primary system. Remote authentication tickets or credentials are sent over the data centers.

10.1.2.13.7 Invisible Takeover

During an invisible takeover the connections between the client and the primary system are kept.

During a takeover you switch your active system from the current primary system to the secondary system.

After an invisible takeover, the client keeps the connections to the primary system and the sessions are restored to the secondary system. This is different from a standard takeover. After a standard takeover, the primary system loses all connections to the client and the secondary system is not aware of the previous connections, which existed between the client and the primary system.

The cross-layer between the session and the client library makes an invisible takeover possible. This cross-layer, called transparent session recovery, recovers the current session state and the physical connection during a takeover or restart.

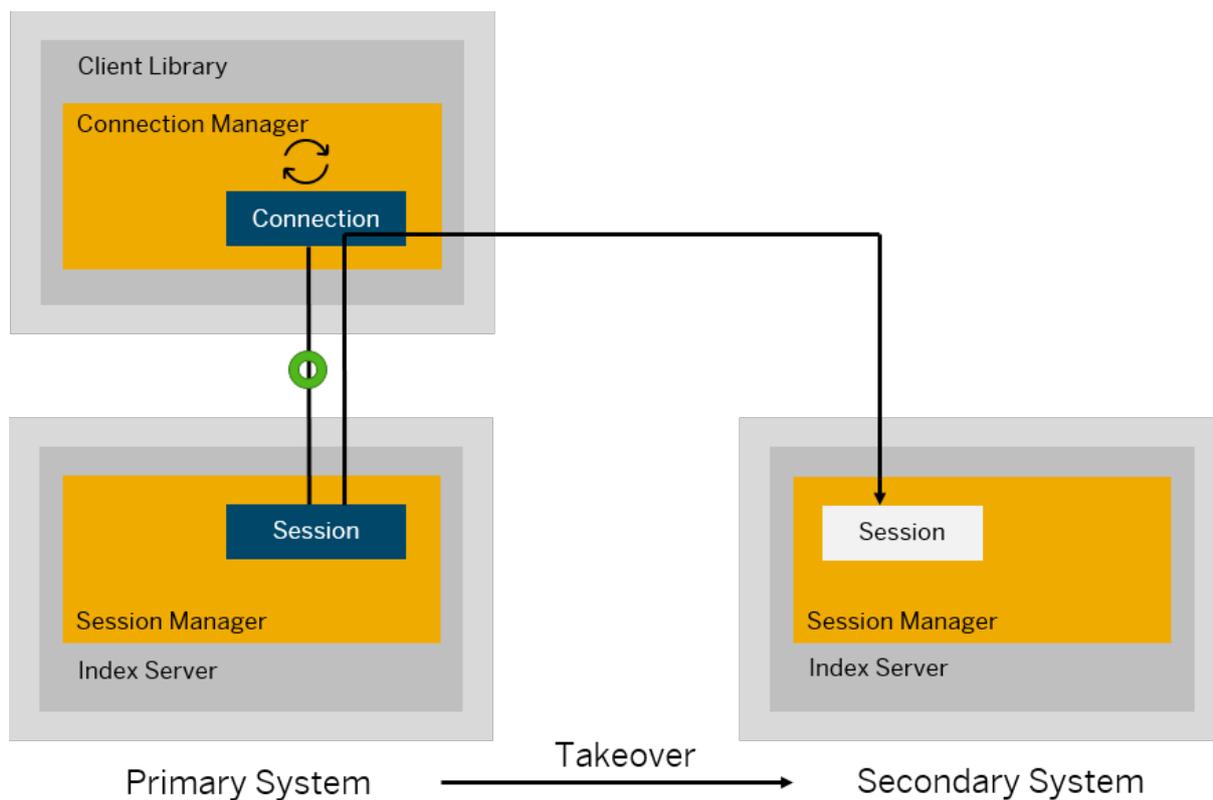
i Note

The transparent session recovery is supported by SQLDBC for SAP HANA 2.0.

An invisible takeover has two functions:

- Keeping the physical connections between the client and the primary and secondary systems
- Restoring the sessions to the secondary system
The session state needs to be recovered and restored to the new primary system. The cross-layer between the session and the client library recovers the current session state and the physical connection into the secondary system.

Use the graphic below for a visualization of an invisible takeover:



Configuration

The `enable_session_recovery` parameter controls the session recovery. When set to `true`, the session variables are recovered and the client connections are restored from the primary system to the secondary system. This parameter is configurable online, but the changes can be applied only to the connections established after making the changes.

Limitations

- All session variables from the current session context should be recovered. However, sessions for which the client executes queries changing their internal variables won't be recovered.
- Sessions which have created or updated either a local or global temporary table with any DDL or DML commands won't be recovered.
- Only read transactions are supported. When the ongoing transaction is a write transaction, the session won't be recovered.

10.1.2.13.8 Virtual IP Address Handling

In an Active/Active (read enabled) configuration, a second virtual IP address for the read access on the secondary system is needed.

Since in an Active/Active (read enabled) configuration both systems are open for SQL access, a second virtual IP address for read access on the secondary system is needed.

During takeover you can keep the virtual IP address of the secondary system. This virtual IP address will be used for read access until a reconnect occurs. The former virtual IP address of the primary system is also rebound to access the former secondary system, which is the now active system. In this situation two virtual IP addresses are available for accessing the former secondary system after takeover. For more information, see *Client Connection Recovery after Takeover*.

i Note

Make sure that the now active system is capable to handle the workload of the former primary system and the read access secondary system.

During failback the system replication systems switch their roles and the virtual IP addresses switch their locations too.

Related Information

[Client Connection Recovery after Takeover \[page 1132\]](#)

10.1.2.13.9 Monitoring Active/Active (Read Enabled)

You can monitor the Active/Active (read enabled) solution using proxy views or the SAP HANA Cockpit.

Proxy views

The embedded statistics server runs in the primary system and collects data from the secondary system providing them in the corresponding proxy schema. For more information, see *Monitoring Secondary Sites*.

SAP HANA Cockpit

The monitoring functionality in the SAP HANA cockpit supports access to the read enabled system. Additionally, it provides information about the delay of the currently available consistent view. For more information, see *Monitoring SAP HANA System Replication with the SAP HANA Cockpit*.

Related Information

[Monitoring Secondary Sites \[page 1194\]](#)

[Monitoring SAP HANA System Replication with the SAP HANA Cockpit \[page 1184\]](#)

10.1.2.14 SAP HANA System Replication with Tenant Databases

The usual SAP HANA system replication principles apply for tenant database systems.

Before you begin preparing a replication strategy for an SAP HANA system, you should be aware of the following important points.

- SAP HANA systems can only be replicated as the whole system. This means that the system database and all tenant databases are part of the system replication. A takeover can only be performed as a whole system. A takeover on the level of a single tenant database is not possible.
- If a new tenant database is created in a configured SAP HANA system replication, it must be backed up to participate in the replication. Afterwards, the initial data shipping is started automatically for this tenant database. If a takeover is done while the initial data shipping is running and not finished, this new tenant database will not be operational after takeover and will have to be recovered with backup and recovery (see the *SAP HANA Database Backup and Recovery* section of the SAP HANA Administration Guide).
- If an active tenant database is stopped in a running SAP HANA system replication, it is stopped on the secondary site as well. If a takeover is done while tenant databases (which were part of the system replication) are stopped, they will be in the same state after takeover as they were on the primary site when they were stopped. The tenant databases must be started to complete the takeover.
- If SAP HANA system replication runs in replication mode SYNC with the full sync option enabled, and if the connection to the secondary site is interrupted, no write operations on the primary site are possible. The operation of creating a tenant database, for example, will wait until the connection to the secondary is reestablished or the SQL statement times out.
- With SAP HANA systems, the services needed are generated automatically in the tenant databases.
- For SAP HANA system replication, a port offset value of 100 is configured to reserve ports for system replication communication. The port number of the replication port is calculated by adding the value for this replication port offset to the internal port number of the corresponding service. Thus, although the same `<instance number>` is used for primary and secondary systems, the `<instance number>+1` is reserved for both systems, because this port range is needed for system replication communication. For SAP HANA systems, this port offset is set to 10000 shifting the ports from the `3<instance number>00` to the `4<instance number>00` port range for the services. This is necessary in SAP HANA system replication with SAP HANA systems, because after `3<instance number>99` is reached new tenant databases allocate port numbers of the next higher instance number.

i Note

To avoid interference with ephemeral ports it might be necessary to adjust the OS port range when using SAP HANA system replication in combination with SAP HANA tenant databases. On Linux this can be accomplished with the following command in the system startup script: `echo "50000 65535" > /proc/sys/net/ipv4/ip_local_port_range.`

- It is possible to copy or move tenant databases between SAP HANA systems. However, you can only use this feature if system replication is not enabled for high availability purposes on either the source or target system for the entire duration of the copy or move process. For more information, see *Copying and Moving Tenant Databases Between Systems*.

i Note

When copying or moving a tenant to the primary system of a system replication landscape, the data shipping for this tenant starts immediately. While this initial data shipping is running, a takeover will cause a loss of data for this tenant on the new primary.

- For SAP HANA systems running with the HIGH isolation level, the system PKI SSFS data and key file must be copied from the primary system to the same location on the secondary system(s). For more information, see *Increase the System Isolation Level* in the SAP HANA Administration Guide.

For more information on the individual points, see the *Availability and Scalability* section of the *SAP HANA Administration Guide*.

Related Information

[Availability and Scalability \[page 1080\]](#)

[SAP HANA Database Backup and Recovery \[page 1229\]](#)

[Copying and Moving Tenant Databases Between Systems \[page 1004\]](#)

[Increase the System Isolation Level \[page 202\]](#)

10.1.2.15 SAP HANA Multitier System Replication

To offer higher levels of availability you can link together multiple systems in a SAP HANA multitier system replication landscape.

You can set up system replication to support geo-clustering, that is multitier system replication between a primary data center and other geographically remote data centers to form a single highly available system.

Related Information

[Setting Up SAP HANA Multitier System Replication \[page 1171\]](#)

[Performing a Failback in SAP HANA Multitier System Replication \[page 1178\]](#)

10.1.2.15.1 Setting Up SAP HANA Multitier System Replication

To offer higher levels of availability you can link together multiple systems in a SAP HANA multitier system replication landscape.

With Multitier System Replication, a tier 2 system replication setup can be used as the source for replication in a chained setup of primary site, tier 2 secondary site and tier 3 secondary site.

After setting up a basic system replication scenario you add a third system to provide another level of redundancy. In a multitier setup the primary system is always on tier 1, a tier 2 secondary has a primary system as its replication source and a tier 3 secondary has the tier 2 secondary as its replication source.

Multitier system replication does not allow operation mode mixtures. However, there is one exception: if the operation mode `logreplay_readaccess` is configured between tier 1 and tier 2, the operation mode `logreplay` can be configured between tier 2 and tier 3. For more information, see *Operation Modes for SAP HANA System Replication*.

Multitier system replication supports various replication mode combinations. For more information, see *Supported Replication Modes between Sites*.

You can configure multitier system replication using the following tools:

- SAP HANA studio
For more information, see *Set Up SAP HANA Multitier System Replication with the SAP HANA Studio*.
- `hdbnsutil`
For more information, see *Set Up SAP HANA Multitier System Replication with `hdbnsutil`*.

Related Information

[Supported Replication Modes between Sites \[page 1172\]](#)

[Set Up SAP HANA Multitier System Replication with `hdbnsutil` \[page 1177\]](#)

[Set Up SAP HANA Multitier System Replication with SAP HANA Studio \[page 1175\]](#)

[Operation Modes for SAP HANA System Replication \[page 1094\]](#)

10.1.2.15.1.1 Supported Replication Modes between Sites

In a multitier system replication scenario, the following replication mode combinations are supported.

Replication Mode Combinations

Tier1 to Tier 2	Tier 2 to Tier 3	Description	Use Case
SYNC	SYNC	<p>In this setup tier 1, tier 2, and tier 3 are coupled with SYNC replication mode.</p> <p>Tier 2 sends the acknowledgment to tier 1 after the log buffer has been received and written to disk, and after the log buffer has also been received and written by tier 3.</p> <p>When primary has received the acknowledge, the buffer has been persisted by all the tiers.</p>	<p>Tier 1 and tier 2 are located in a local data center for fast takeover.</p> <p>Tier 3 is used for disaster recovery in a second close-by data center.</p>
SYNC	SYNCMEM	<p>Tier 2 sends the acknowledge to tier 1 after the log buffer has been received, written to disk and it has been also received by tier 3.</p> <p>When the primary receives acknowledgment, it is not clear that also tier 3 has persisted the buffer to disk, but disk IO on tier 3 has been triggered.</p>	<p>Tier 1 and tier 2 are located in a local data center for fast takeover.</p> <p>Tier 3 is used for disaster recovery in a second close-by data center.</p>
SYNC	ASync	<p>Tier 1 and tier 2 are closely coupled with replication mode SYNC, while tier 3 is decoupled by using ASync.</p>	<p>Tier 1 and tier 2 are located in a local data center for fast takeover.</p>

Tier1 to Tier 2	Tier 2 to Tier 3	Description	Use Case
		<p>Tier 2 acknowledges the arrival of the redo log buffers in-memory and on disk to tier 1, while it only hands over the redo log buffer to the network without awaiting an acknowledgment from tier 3.</p> <p>If the connection to tier 3 is too slow and the ASYNC replication buffer (an intermediate memory buffer) is running full, ASYNC replication to tier 3 can have an impact on the primary.</p>	<p>Tier 3 is used for disaster recovery in a far distant data center.</p>
SYNCMEM	SYNC	<p>In this synchronous setup tier 1 and tier 2 are closely coupled with replication mode SYNCMEM, while tier 3 is closely coupled with SYNC.</p> <p>Tier 2 sends the acknowledgment to tier 1 after the log buffer has been received in memory. IO is triggered asynchronously. The asynchronous IO also triggers the send operation to tier 3. The log write on tier 2 is confirmed, when also tier 3 has written the log buffer.</p> <p>When the primary receives the acknowledge, it is unclear, if tier 3 has already received and persisted the log buffer.</p>	<p>Tier 1 and tier 2 are located in a local data center for fast takeover.</p> <p>Tier 3 is used for disaster recovery in a second close-by data center.</p>

Tier1 to Tier 2	Tier 2 to Tier 3	Description	Use Case
SYNCMEM	SYNCMEM	<p>In this setup tier 1, tier 2, and tier 3 are coupled with replication mode SYNCMEM.</p> <p>Tier 2 sends the acknowledgment to tier 1 after the log buffer has been received in memory. IO is triggered asynchronously. The asynchronous IO also triggers the send operation to tier 3. The log write on tier 2 is confirmed, when tier 3 has received the log buffer in memory.</p> <p>When the primary receives the acknowledge, it is unclear, if tier 3 has already received and persisted the log buffer.</p>	<p>Tier 1 and tier 2 are located in a local data center for fast takeover.</p> <p>Tier 3 is used for disaster recovery in a second close-by data center.</p>
SYNCMEM	ASync	<p>Tier 1 and tier 2 are closely coupled with replication mode SYNCMEM, while tier-3 is decoupled with ASync replication.</p> <p>Tier 2 acknowledges the arrival of the redo log buffers in-memory to tier 1, while it only hands over the redo log buffer to the network without awaiting an acknowledgment from tier 3.</p> <p>If the connection to tier 3 is too slow and the ASync replication buffer (an intermediate memory buffer) is running full, ASync</p>	<p>Primary and tier 2 are used in a local data center for fast takeover.</p> <p>Tier 3 is used for disaster recovery in a far distant data center.</p>

Tier1 to Tier 2	Tier 2 to Tier 3	Description	Use Case
		replication can have an impact on the primary.	
ASYN	ASYN	<p>With these asynchronous replication modes there is no wait for acknowledgments between tiers (no acknowledge propagation).</p> <p>A replication backlog for tier 2 and tier 3 is possible.</p> <p>Information about the replication status on tier 1 and tier 2 is available in the ASYN replication buffer (an intermediate memory buffer). This buffer running full could cause a minimal impact on the performance of the primary.</p>	<p>Tier 1 performance is most important as well as a disaster recovery capability. For best performance of tier 1 decouple tier 2 and tier 3.</p> <p>Data loss on tier 2 and tier 3 is possible to some extent, but performance is more critical.</p>

10.1.2.15.1.2 Set Up SAP HANA Multitier System Replication with SAP HANA Studio

You can set up SAP HANA multitier system replication using the SAP HANA studio.

Prerequisites

- You have considered all the general prerequisites needed to set up system replication. For more information, see *General Prerequisites for Setting Up SAP HANA System Replication*.
- You have installed and configured three identical, independently operational SAP HANA systems – a primary system, a tier 2 secondary system and a tier 3 secondary system.
- You have added the systems in the SAP HANA studio.
- You have verified that the `log_mode` parameter in the `persistence` section of the `global.ini` file is set to **normal** for the systems.
You can do this in the Administration editor (*Configuration* tab) of the SAP HANA studio.

- You have performed a data backup on the tier 2 secondary system.
- You have stopped the tier 3 secondary system.

Context

The following procedure describes how to add a tier 3 secondary with a synchronously running tier 2 system replication.

Procedure

1. Enable system replication on the tier 2 secondary, which has to be online, as follows:
 - a. In the *Systems* view right click the tier 2 secondary system, choose ► *Configuration and Monitoring* ► *Configure System Replication* ▾
The *Configure System Replication* dialog opens. The *Enable System Replication* option is selected by default. The site name is already known from the topology metadata.
 - b. Choose *Next*.
 - c. Review the configured information and choose *Finish*.
2. Register the tier 3 secondary system as follows:
 - a. Stop the tier 3 secondary system if it is still running. Right-click the tier 3 secondary system and choose ► *Configuration and Monitoring* ► *Stop System* ▾
 - b. In the *Systems* view, right-click the tier 3 secondary system and choose ► *Configuration and Monitoring* ► *Configure System Replication* ▾.
The *Configure System Replication* dialog opens.
 - c. Choose *Register Secondary System* and then *Next*.
 - d. Enter the required system information and the logical name used to represent the tier 3 secondary system.

i Note

If you are operating a distributed system on multiple hosts, you enter the name of the host on which the master name server is running.

 - e. Specify the log replication mode *Asynchronous (mode=async)* and enter the tier 2 secondary system's host name:
 - f. Review the configured information and choose *Finish*.

Results

The secondary system is automatically started and the replication process to the tier 3 secondary then starts automatically.

Related Information

[General Prerequisites for Configuring SAP HANA System Replication \[page 1090\]](#)

[Supported Replication Modes between Sites \[page 1172\]](#)

10.1.2.15.1.3 Set Up SAP HANA Multitier System Replication with hdbnsutil

You can set up SAP HANA multitier system replication with hdbnsutil.

Prerequisites

- You have considered all the general prerequisites needed to set up system replication. For more information, see *General Prerequisites for Setting Up SAP HANA System Replication*.
- You have installed and configured three identical, independently operational SAP HANA systems – a primary system, a tier 2 secondary system and a tier 3 secondary system.

Context

The following steps show how to set up such a system. In this scenario there are three SAP HANA systems: A, B and C, named SiteA, SiteB and SiteC. Furthermore, in this scenario multitier system replication supports a tier 2 secondary with sync replication mode and a tier 3 secondary with async replication mode.

Procedure

1. [A] Start the SAP HANA database.
2. [A] Create a data backup or storage snapshot. In multiple-container systems, the system database and all tenant databases must be backed up.
3. [A] Enable system replication and give the system a logical name. As `<sid>adm`:

```
cd /usr/sap/<sid>/HDB<instancenr>/exe
```

```
./hdbnsutil -sr_enable --name=SiteA
```

4. Stop the tier 2 secondary.

As `<sid>adm` run the SAPControl program to shut down the system:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StopSystem  
HDB
```

5. [B] On the stopped tier 2 secondary, register site B with Site A as <sid>adm:

```
hdbnsutil -sr_register --replicationMode=sync --name=SiteB
--remoteInstance=<instId> --remoteHost=<hostname_of_A>
```

6. [B] Start the tier 2 secondary system.

As <sid>adm run the SAPControl program to start the system:

```
/usr/sap/hostctrl/exe/sapcontrol sapcontrol -nr <system number> -function
StartSystem HDB
```

7. [B] Enable this site as the source for a tier 3 secondary system:

As <sid>adm on the tier 2 secondary run `hdbnsutil -sr_enable`

8. [C] Stop the tier 3 secondary system.

As <sid>adm run the SAPControl program to shut down the system:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StopSystem
HDB
```

9. [C] On the stopped system, register siteC as a tier 3 secondary system as <sid>adm:

```
hdbnsutil -sr_register --replicationMode=async --name=SiteC
--remoteInstance=<instId> --remoteHost=<hostname_of_B>
```

10. [C] Start the SAP HANA database on the tier 3 secondary.

As <sid>adm run the SAPControl program to start the system:

```
/usr/sap/hostctrl/exe/sapcontrol sapcontrol -nr <system number> -function
StartSystem HDB
```

11. Check the replication status in the SAP HANA studio ► *landscape* ► *replication* ► tab or with the M_SERVICE_REPLICATION system view.

Related Information

[General Prerequisites for Configuring SAP HANA System Replication \[page 1090\]](#)

[Supported Replication Modes between Sites \[page 1172\]](#)

[Operation Modes for SAP HANA System Replication \[page 1094\]](#)

10.1.2.15.2 Performing a Failback in SAP HANA Multitier System Replication

If the primary system failed, a takeover to the tier 2 secondary system was done. Once your failed site is operational again you can attach it as a tier 3 secondary system or you can restore the original multitier system replication configuration.

For more information, see *Attach the Original Primary System as a New Tier 3 Secondary System* and *Restore the Original SAP HANA Multitier System Replication Configuration*.

Related Information

[Attach the Original Primary System as a New Tier 3 Secondary System \[page 1179\]](#)

[Restore the Original SAP HANA Multitier System Replication Configuration \[page 1180\]](#)

10.1.2.15.2.1 Attach the Original Primary System as a New Tier 3 Secondary System

Once your failed site is operational again you can attach it as a tier 3 secondary system.

Context

The steps below show how to set up multitier system replication again after a takeover. In these scenarios there are three SAP HANA systems A, B and C, named SiteA, SiteB, and SiteC. Furthermore, in this scenario multitier system replication supports a tier 2 secondary with sync replication mode and a tier 3 secondary with async replication mode.

Procedure

SiteA failed, SiteB has taken over and now you attach SiteA as the tier 3 secondary.

1. [C] Change the replication mode of the new tier 2 secondary:

```
cd /usr/sap/<sid>/HDB<instance_number>/exe
./hdbnsutil -sr_changemode --replicationMode=sync
```

Multitier system replication supports various replication mode combinations. For more information, see *Supported Replication Modes between Sites*.

2. [C] Enable SiteC as the replication source:

```
hdbnsutil -sr_enable
```

3. [A] Make sure that the SAP HANA database is stopped. This should be the case as a takeover was already carried out otherwise you can stop it with the following command:

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function StopSystem
HDB
```

4. [A] Register SiteA as a new tier 3 secondary.

```
hdbnsutil -sr_register --replicationMode=async --name=SiteA --
remoteInstance=<instId> --remoteHost=<hostname_of_C>
```

5. [A] Start the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function StartSystem  
HDB
```

6. [B] Check in M_SERVICE_REPLICATION that sync system replication is ACTIVE from SiteB to SiteC and that async replication is ACTIVE from SiteC to SiteA.

Related Information

[Supported Replication Modes between Sites \[page 1172\]](#)

10.1.2.15.2.2 Restore the Original SAP HANA Multitier System Replication Configuration

Once your failed site is operational again you can restore the original SAP HANA multitier system replication configuration.

Context

The steps below show how to set up multitier system replication again after a takeover. In these scenarios there are three SAP HANA systems A, B and C, named SiteA, SiteB, and SiteC. Furthermore, in this scenario multitier system replication supports a tier 2 secondary with sync replication mode and a tier 3 secondary with async replication mode.

Procedure

You want to restore the original multitier setup:

1. [C] Stop the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function StopSystem  
HDB
```

2. [C] Unregister SiteC from SiteB:

```
hdbnsutil -sr_unregister
```

3. [A] Stop the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function StopSystem  
HDB
```

4. [A] Register as secondary:

```
hdbnsutil -sr_register --replicationMode=sync --name=SiteA --
remoteInstance=<instId> --remoteHost=<hostname_of_B>
```

5. [A] Start the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function StartSystem
HDB
```

6. [B] Check in M_SERVICE_REPLICATION that sync system replication is ACTIVE from SiteB to SiteA.

7. [A] SiteA takes over as the primary system:

```
hdbnsutil -sr_takeover
```

8. [B] Stop the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function StopSystem
HDB
```

9. [A] Enable system replication:

```
hdbnsutil -sr_enable --name=SiteA
```

10. [B] Register SiteB as the tier 2 secondary of SiteA.

```
hdbnsutil -sr_register --replicationMode=sync --name=SiteB --
remoteInstance=<instId> --remoteHost=<hostname_of_A>
```

11. [B] Start the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function StartSystem
HDB
```

12. [B] Enable SiteB as a replication source system:

```
hdbnsutil -sr_enable
```

13. [C] Register SiteC as a tier 3 secondary in the multitier system replication scenario:

```
hdbnsutil -sr_register --replicationMode=async --name=SiteC --
remoteInstance=<instId> --remoteHost=<hostname_of_B>
```

14. [C] Start the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function StartSystem
HDB
```

15. [B] Check in M_SERVICE_REPLICATION that sync replication is ACTIVE from SiteA to SiteB and that async replication is ACTIVE from SiteB to SiteC.

10.1.2.16 Monitoring SAP HANA System Replication

To ensure rapid takeover in the event of planned or unplanned downtime, you can monitor the status of replication between the primary system and the secondary system.

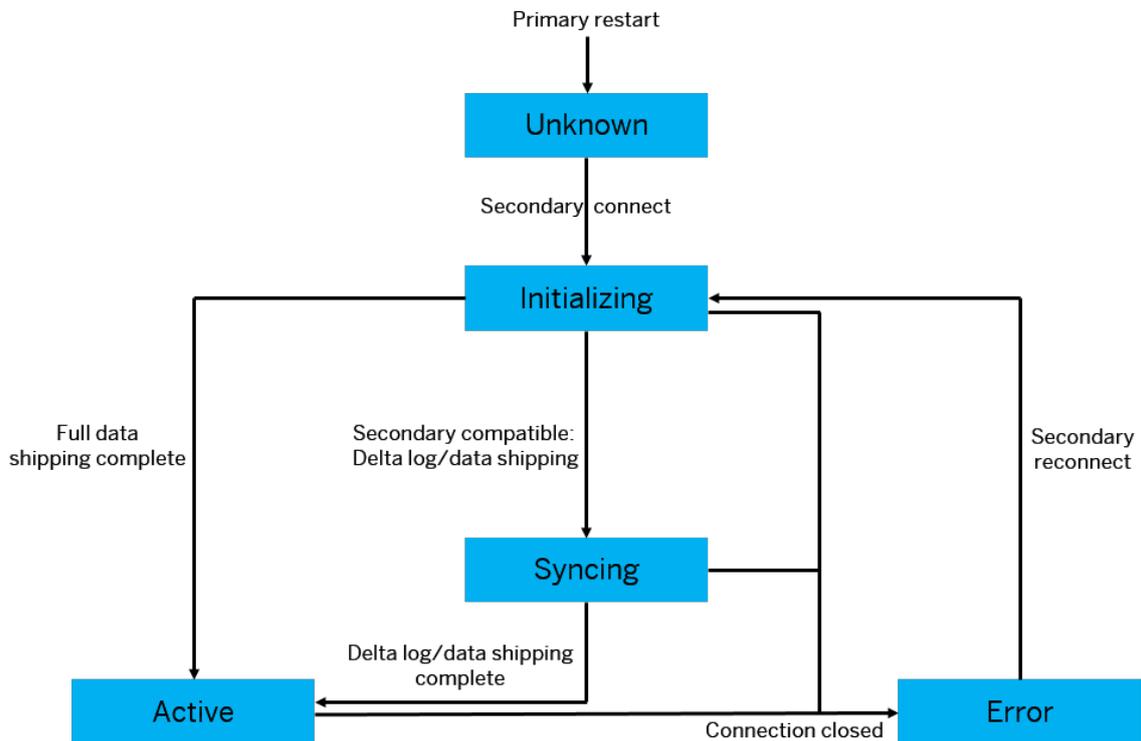
You can monitor system replication using the following tools:

- SAP HANA cockpit
For more information, see *Monitoring System Replication with the SAP HANA Cockpit*.
- SAP HANA studio
For more information, see *Monitoring System Replication with the SAP HANA Studio*.
- hdbnsutil
For more information, see *Monitoring System Replication with hdbnsutil*.

Replication Status

The current status of system replication can be checked with all of the mentioned tools:

Status	Description
UNKNOWN	Secondary did not connect to primary since last restart of the primary
INITIALIZING	Initial data transfer in progress. In this state, the secondary is not usable at all.
SYNCING	Secondary is syncing again (for example, after a temporary connection loss or restart of the secondary).
ACTIVE	Initialization or sync with primary is complete and secondary is continuously replicating. No data loss will occur in SYNC mode.
ERROR	Error occurred on the connection (additional details on the error can be found in REPLICATION_STATUS_DETAILS).



Configure E-mail Notifications

To receive e-mail notification of alerts, you can configure check 78, 79, and 94 accordingly, for more information see *Configure E-Mail Notifications for Alerts*.

- Alert ID 78: System Replication Connection Closed
- Alert ID 79: System Replication Configuration Parameter Mismatch
- Alert ID 94: Log Replay Backlog for System Replication Secondary

Related Information

[Monitoring SAP HANA System Replication with the SAP HANA Cockpit \[page 1184\]](#)

[Monitoring SAP HANA System Replication with the SAP HANA Studio \[page 1185\]](#)

[Monitoring SAP HANA System Replication with Command Line Script \[page 1186\]](#)

[Monitoring Status and Resource Usage of System Components \[page 364\]](#)

[Monitoring Alerts \[page 373\]](#)

[Configure E-Mail Notifications for Alerts \[page 376\]](#)

[Monitoring and Replicating INI File Parameter Changes \[page 1190\]](#)

[SAP HANA System Replication Alerts \[page 1192\]](#)

[SAP HANA SQL and System Views Reference](#)

10.1.2.16.1 Monitoring SAP HANA System Replication with the SAP HANA Cockpit

To monitor SAP HANA system replication, you can use the [System Replication](#) tile in the SAP HANA cockpit.

To open the [System Replication Overview](#) page, click the [System Replication](#) tile on the [Overview](#) page in the SAP HANA cockpit.

The [System Replication Overview](#) displays a graphical representation of the system replication landscape with the following information:

- The name and role of the system, as well as the selected operation mode
For the operation modes `logreplay` and `logreplay_readaccess` a retention time estimation is also displayed. The [Estimated log retention time](#) is an estimation of the time left before the primary system starts to overwrite the `RetainedFree` marked log segments and a full data shipping becomes necessary to get the primary and secondary systems back in sync after a disconnect situation. The [Estimated log full time](#) is an estimation of the time left before the primary system runs into a log full. The value shown in the header shows the situation into which the system could run first: log retention or log full.
- If the SQL ports of the secondary system are open for read access
- The replication mode used between the systems
- The current average redo log shipping time and the average size of shipped redo log buffers
It describes how long it took on average to send redo log buffers to the secondary site based on measurements of the last 24 hours.

Furthermore, detailed information on system replication is provided in the following four tabs:

i Note

These tabs are displayed only if you configured a system replication before.

Tabs on the System Replication Overview

Tab Name	Description
Related Alerts	The Related Alerts tab provides a description of any existing alerts, as well as their priority. This tab is only displayed when system replication related alerts are available.
Replicated Services	The Replicated Services tab provides information on the replication status per site and service.
Network	The Network tab provides information on the time it took to ship the redo log to the secondary system and to write the redo log to the local log volume on disk. You can select the network connection that you want to analyze (for example, Network Site 1 to 2 or Network Site 2 to 3). The graph displayed compares the local write wait time with the remote write wait time monitored over the last 24 hours.
Log Shipping Backlog	The Log Shipping Backlog tab provides a graphical representation on the history of the log shipping backlog.

Tab Name	Description
Log Replay	<p>The Log Replay tab provides a graphical representation on the delay of the secondary system. This tab is displayed if the chosen operation mode for the system replication landscape is <code>logreplay</code> or <code>logreplay_readaccess</code>.</p> <p>When this tab is activated for a secondary system, the log replay delay is shown for the last 24 hours.</p> <p>Furthermore, in this tab you can select to visualize the estimated log retention time as well as the estimated log full time for all system replication relevant services.</p>
Network Speed Check	<p>The Network Speed Check tab provides a way to measure the network speed of the system replication host-to-host network channel mappings.</p>
Network Security Settings	<p>The Network Security Settings tab displays the specific network security details configured between the primary and the secondary systems.</p>

10.1.2.16.2 Monitoring SAP HANA System Replication with the SAP HANA Studio

You can monitor SAP HANA system replication using the SAP HANA studio.

You can monitor system replication in the Administration editor of the primary system as follows:

- The general status is displayed on the [Overview](#) tab.
- Detailed information is available on the [Landscape > System Replication](#) tab. Here you can see the system replication status. For more information on the system replication status, see [Monitoring System Replication](#).
Since the secondary instance does not accept SQL connections while data replication is active, basic information about the secondary system is also shown.
For more information about the meaning of the individual fields, see the system view `M_SERVICE_REPLICATION`.

Related Information

[Monitoring SAP HANA System Replication \[page 1181\]](#)

10.1.2.16.3 Monitoring SAP HANA System Replication with Command Line Script

You can monitor SAP HANA system replication using a command line script.

Check the overall status of the system replication using as <sid>adm OS user the script systemReplicationStatus.py (located in \$DIR_INSTANCE/ /exe/python_support).

```
ld2131:- # su - utladm
utladm@ld2131:/usr/sap/UT1/HDB01> python exe/python_support/systemReplicationStatus.py
| Host | Port | Service Name | Volume ID | Site ID | Site Name | Secondary | Secondary | Secondary | Secondary | Secondary | Replication | Replication | Replication |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| ld2131 | 30107 | xsengine | 2 | 1 | SiteA | ld2132 | 30107 | 2 | SiteB | YES | SYNC | ACTIVE |
| ld2131 | 30101 | nameserver | 1 | 1 | SiteA | ld2132 | 30101 | 2 | SiteB | YES | SYNC | ACTIVE |
| ld2131 | 30103 | indexserver | 3 | 1 | SiteA | ld2132 | 30103 | 2 | SiteB | YES | SYNC | ACTIVE |
| ld2132 | 30107 | xsengine | 2 | 2 | SiteB | ld2133 | 30107 | 3 | SiteC | YES | ASYNC | ACTIVE |
| ld2132 | 30101 | nameserver | 1 | 2 | SiteB | ld2133 | 30101 | 3 | SiteC | YES | ASYNC | ACTIVE |
| ld2132 | 30103 | indexserver | 3 | 2 | SiteB | ld2133 | 30103 | 3 | SiteC | YES | ASYNC | ACTIVE |

status system replication site "2": ACTIVE
status system replication site "3": ACTIVE
overall system replication status: ACTIVE

Local System Replication State
-----
mode: PRIMARY
site id: 1
site name: SiteA
```

The script provides the following return codes:

- 10: No System Replication
- 11: Error
- 12: Unknown
- 13: Initializing
- 14: Syncing
- 15: Active

Related Information

[Monitoring Status and Resource Usage of System Components \[page 364\]](#)

[Monitoring Alerts \[page 373\]](#)

[Configure E-Mail Notifications for Alerts \[page 376\]](#)

[Monitoring and Replicating INI File Parameter Changes \[page 1190\]](#)

[SAP HANA System Replication Alerts \[page 1192\]](#)

[SAP HANA SQL and System Views Reference](#)

10.1.2.16.4 SAP HANA System Replication Details

Detailed information about system replication.

General Overview

Column	Description
Site ID 1	Generated ID of the primary site

Column	Description
Secondary Site ID 2	Generated ID of the secondary site
Service	Name of the service
Volume ID	Persistence volume ID
Operation Mode	<ul style="list-style-type: none"> LOGREPLAY DELTA DATA SHIPPING LOGREPLAY_READACCESS
Replication Mode	<p>Configured replication mode:</p> <ul style="list-style-type: none"> SYNC: synchronous replication with acknowledgement when buffer has been written to disk SYNCMEM: synchronous replication with acknowledgement when buffer arrived in memory ASYNC: asynchronous replication UNKNOWN: is set if replication mode could not be determined (this might be the case, for example, if there are communication errors when getting status information from a service).
Replication Status	<p>Current status of replication:</p> <ul style="list-style-type: none"> UNKNOWN: secondary did not connect to primary since last restart of the primary INITIALIZING: initial data transfer occurs, in this state, the secondary is not usable at all SYNCING: secondary is syncing again (for example, after a temporary connection loss or restart of the secondary) ACTIVE: initialization or sync with primary is complete and secondary is continuously replicating. If crash occurs, no data loss will occur in SYNC mode. ERROR: error occurred with the connection (details can be found in Replication Details)
Replication Details	Additional information for Replication Status, for example, the error text if status is ERROR.

Column	Description
Full Sync	<p>Indicates if the service is currently operating in full sync.</p> <p>If full sync is enabled (global.ini/[system_replication]/enable_full_sync) in a running system, full sync might not be active immediately. This is done to prevent the system from blocking transactions immediately when setting the parameter to true. Instead, in a first step, full sync has to be enabled by the administrator. In a second step it is internally activated, when the secondary is connected and becomes ACTIVE.</p> <ul style="list-style-type: none"> • DISABLED: full sync is not configured at all (global.ini/[system_replication]/enable_full_sync = false) • ENABLED: full sync is configured, but it is not yet active, so transactions do not block in this state. To become active the secondary has to connect and Replication Status has to be ACTIVE. • ACTIVE: full sync mode is configured and active. If a connection of a connected secondary is getting closed, transactions on the primary side will block in this state. <p>If full sync is enabled when an active secondary is currently connected, the FULL_SYNC will be immediately set to ACTIVE.</p>
Secondary Fully Recoverable	<p>TRUE: No full data backup is needed after takeover on secondary. Backups created on the primary and local log segments enable a full database recovery.</p> <p>FALSE: Log segments needed for a full database recovery are missing. After takeover a full data backup has to be executed before a full recovery up to the most recent time point can be executed.</p>
Secondary Active	Status of the secondary node (also see ACTIVE_STATUS in M_SERVICES)
Secondary Connect Time	Timestamp the secondary connected to the primary. If there are reconnects from the secondary side, this field contains the last connect time.
Secondary Reconnect Count	Number of reconnects from secondary side for this service.
Secondary Failover Count	Number of failovers for this service on secondary side.
Buffer Full count	Number of times, the asynchronous replication buffer was full since last service restart (0 for replication modes sync/syncmem).

Log Positions

Column	Description
Last Log Position	Last known log position on primary
Last Log Position Time	Timestamp of last known log position

Column	Description
Replayed Log Position	Log end position of the last known replayed log buffer on secondary site
Replayed Log Position Time	Timestamp of the last known replayed log buffer on the secondary site
Last Shipped Log Position Time	Timestamp of last log position being shipped to secondary
Shipped Log Buffer Count	Number of log buffers shipped to secondary
Shipped Log Buffers Total Size (Bytes)	Size of all log buffers shipped to secondary
Shipped Log Buffers Total Time (µs)	Time taken to ship all the log buffers to the secondary. <ul style="list-style-type: none"> • SYNC/SYNCMEM: total round trip time to send the log buffers and receive the acknowledgment. • ASYNC: start time when sending the log buffers, end time when the OS reports that the log buffers were sent (and the log shipping buffer space was freed). This could be shorter than the SYNC/SYNCMEM duration
Time delay (ms)	Size delay between the last shipped log position time and the replayed log position time on the secondary
Size delay (Bytes)	Time delay between the last shipped log position size and the replayed log position size on the secondary (1 log position = 64 bytes)

Savepoints

Column	Description
Last Savepoint Version	Last savepoint version on primary
Last Savepoint Log Position	Log position of current savepoint
Last Savepoint Start Time	Timestamp of current savepoint
Last Shipped Savepoint Version	Last savepoint version shipped to secondary
Last Shipped Savepoint Log Position	Log position of last shipped savepoint
Last Shipped Savepoint Time	Timestamp of last shipped savepoint

Full Data Replica

Column	Description
Full Data Replica Shipped Count	Number of full data replicas shipped to secondary
Full Data Replica Shipped Total Size (Bytes)	Total size of all full backups shipped to secondary
Full Data Replica Shipping Total Time (µs)	Duration for shipping all full data replica
Last Full Data Replica Shipped Size (Bytes)	Size of last full data replica shipped to secondary
Start Time of Last Full Data Replica	Start time of last full data replica
End Time of Last Full Data Replica	End time of last full data replica

Delta Data Replica

This information is only displayed if the operation mode is `logreplay`.

Column	Description
Delta Data Replica Shipped Count	Number of delta data replicas shipped to secondary
Delta Data Replica Shipped Total Size (Bytes)	Total size of all delta data replicas shipped to secondary
Delta Data Replica Shipped Total Time (µs)	Duration for shipping of all delta data replicas
Size of Last Delta Data Replica (Bytes)	Size of last delta data replica
Start Time of Last Delta Data Replica	Start time of last data delta replica
End Time of Last Delta Data Replica	End time of last data delta replica

Backlog

This information is only displayed if the operation mode is `logreplay`.

Column	Description
Current Replication Backlog Size (Bytes)	<p>Current replication backlog in bytes, this means, size of all log buffers that have been created on primary site, but not yet sent to the secondary site.</p> <p>Even in replication modes <code>sync/syncmem</code> this column can have a value different from 0.</p> <p>Here it represents the size of log buffers that are in the local send queue (max number of those buffers is the number configured log buffers on primary site).</p>
Max Replication Backlog Size (Bytes)	Max replication backlog in bytes (max value of <code>BACKLOG_SIZE</code> since system start).
Current Replication Backlog Time (µs)	<p>Current replication backlog in microseconds. This is time difference between time of the last sent log buffer and the current log buffer.</p> <p>Even in replication modes <code>sync/syncmem</code> this column can have a value different from 0, because log buffers are still in the send queue (max number of these buffers is the number of log buffers configured on primary site).</p>
Max Replication Backlog Time (µs)	Max replication backlog in microseconds (max value of <code>BACKLOG_TIME</code> since system startup).

10.1.2.16.5 Monitoring and Replicating INI File Parameter Changes

To check, if the `.ini` file parameters are the same on each site of a system replication landscape the configuration parameter checker reports on any differences between primary, secondary, and tier 3 secondary systems.

Some parameters may have different settings on the primary and the secondary sites on purpose; one example is the `global_allocation_limit` parameter where the secondary is used for other systems. By adding those parameters to the below exclusion list they are excluded from checking and replication.

With parameter replication activated, any changes made on the primary are automatically replicated to the secondary sites; without this parameter replication activated changes should be manually duplicated on the other system.

In the current version of the configuration parameter checker, the checks:

- Are done every hour by default
- Generate alerts, visible both in SAP HANA studio and the system view M_EVENTS.
- Are optimized for the most recently changed parameters.

Enable and disable the parameter check on the primary site with `[inifile_checker]/enable = true | false`

The parameter checker is on by default.

Enable and disable the parameter replication on the primary site with `[inifile_checker]/replicate = true | false`

The parameter replication is off by default.

You can replicate the .ini file parameters based on the alerts as follows:

Parameter on the Primary System	Parameter on the Secondary System	Activity
set	not set	Copy parameter to the secondary system.
not set	set	Delete parameter on the secondary system.
set to value x	set to value y	Copy value x to the secondary system.

The parameter changes on the secondary system are applied differently for each parameter:

- Online changeable parameters become active after the `ALTER SYSTEM` command or by editing the .ini file followed by the automatically triggered `hdbnsutil -reconfig` command.
- Offline changeable parameters become active after a restart. When changing such a parameter it is necessary to restart the primary and secondary systems before a takeover. For more information, see SAP Note 2036111.

To prevent parameters from generating alerts and getting replicated eventually, it is possible to create exclusions. In the following example, different global allocation limits (GAL) on primary and secondary systems can be set without being overwritten by the parameter replication:

```
[inifile_checker]
enable = true|false
interval = 3600
exclusion_global.ini/SYSTEM = memorymanager/global_allocation_limit
```

The exclusion rules are written in the following syntax (comma separated list) and take effect immediately:

```
exclusion_[inifile name|*][<LAYER>] = [section with
wildcards|*][/parameter with wildcards|*], ...
<LAYER> := SYSTEM\|HOST\|DATABASE\|\"
```

Related Information

[Configuring SAP HANA System Properties \(INI Files\) \[page 291\]](#)

[SAP Note 2036111](#)

10.1.2.16.6 SAP HANA System Replication Alerts

Alerts issued by the primary system warn you of potential problems.

The following alerts are issued by the primary system:

Primary System Alerts

Alert	Description
Alert ID 78: System Replication Connection Closed	Alerts 78 and 79 are raised when a system replication connection is closed or when there is a system replication configuration parameter mismatch.
Alert ID 79: System Replication Configuration Parameter Mismatch	<p>Starting with SAP HANA 1.0 SPS 09 these two alerts cover the distinct situations where the connection to the secondary site is closed or where there is a configuration parameter mismatch between the replication sites. These alerts require that you have migrated to the new statistics service (see SAP Note 1917938).</p> <p>Before SAP HANA 1.0 SPS 09 there was one alert, categorized as an "Internal Event" (Alert 21). It was created when:</p> <ul style="list-style-type: none">• The connection to the secondary site was closed.• There was a configuration parameter mismatch between the replication sites. <p>Both situations were covered by one event type and could only be distinguished by the information text provided.</p> <p>Since SAP HANA 1.0 SPS 11 old style alerts based on alert 21 are not created anymore as a default.</p> <p>You can create them by setting the configuration parameter <code>keep_old_style_alert</code> to <code>true</code> in the system replication section of the <code>global.ini</code> file. These alerts can be required to keep the existing monitoring infrastructure, which relies on them, working. If activated, new alerts and old style alerts are created in parallel.</p>

Alert	Description
Alert ID 94: System Replication Logreplay Backlog	<p>Alert 94 is raised when the system replication logreplay backlog is increased. In this case, logreplay is delayed on the secondary site causing a longer takeover time.</p> <p>The alert has a different priority based on the threshold reached:</p> <ul style="list-style-type: none"> • Low: 10 GB < logreplay backlog < 50 GB • Medium: 50 GB <= logreplay backlog < 500 GB • High: logreplay backlog >= 500 GB <p>To identify the reason for the increased system replication logreplay backlog, check the state of the services on the secondary system. To get more information, monitor the secondary site. Possible causes for the increased system replication logreplay backlog can be, for example, a slow or not functioning log replay, or non-running service on the secondary system.</p>
Alert ID 104: System Replication Increased Log Shipping Backlog	<p>Alert 104 is raised when the system replication log shipping backlog is increased. In this case, the log shipping to the secondary system is delayed or will not work properly causing data loss on the secondary system in case a takeover is executed.</p> <p>The alert has a different priority based on the threshold reached:</p> <ul style="list-style-type: none"> • Low: 1 GB < log shipping backlog < 10 GB • Medium: 10 GB <= log shipping backlog < 50 GB • High: log shipping backlog >= 50 GB <p>To identify the reason for the increased system replication log shipping backlog, check the status of the secondary system. Possible causes for the increased system replication log shipping backlog can be a slow network performance, connection problems, or other internal issues (for example, in sync or syncmem replication modes).</p>
Alert ID 106: ASYNC Replication In-Memory Buffer Overflow	<p>Alert 106 is raised when the local in-memory buffer in the ASYNC replication mode is running full indicating possible network issues with the connection to the secondary system.</p> <p>The alert has a different priority based on the threshold reached:</p> <ul style="list-style-type: none"> • Medium: if full once within 24 hours • High: if full more than once within 24 hours <p>To identify the reason for the local in-memory buffer running full, check the buffer size, the network, the IO on the secondary system, or look for peak loads.</p> <p>The alert depends on the setting of the <code>logshipping_async_wait_on_buffer_full</code> and <code>logshipping_async_wait_on_buffer_full</code> parameter. For more information about these parameters, see <i>SAP HANA System Replication Configuration Parameters</i>.</p>

For information on alerts issued by hosts of the secondary system, see *Monitoring Secondary Sites*.

Related Information

[Monitoring Secondary Sites \[page 1194\]](#)

[SAP HANA System Replication Configuration Parameters \[page 1109\]](#)

[SAP Note 1917938](#)

10.1.2.16.7 Monitoring Secondary Sites

Remote SQL access on the primary site allows monitoring and reporting of the secondary site statistics.

Proxy schemas and views are provided on the primary site which extract the corresponding information from the monitoring views on the secondary site. The retrieval of statistics is unaffected by the replication or operation mode and is available for a two system replication setup as well as for multitier landscapes.

Alerts issued by secondary system hosts are displayed in the *Alerts* app of the SAP HANA cockpit.

A new schema is created on the primary site for each registered secondary site. This schema follows the naming convention `_SYS_SR_SITE_<siteName>`, where `<siteName>` is the case-sensitive name given at registration time of the secondary. This schema contains a selected subset of monitoring views (for example, `M_VOLUME_IO_TOTAL_STATISTICS`), which proxies the statistics from the secondary site.

i Note

If system replication is configured as an Active/Active (read enabled) system with the `logreplay_readaccess` operation mode, then in the proxy schema `_SYS_SR_SITE_<siteName>` more data is available from the secondary system. More monitoring views of the secondary system can be accessed via virtual tables.

These proxy views have the same column definitions as the equivalently named public synonyms already available for the primary.

When a secondary site is unregistered the corresponding schema will be dropped.

Limitations

- Monitoring view access is only possible if the primary and secondary site run with exactly the same software version.
- When such a proxy view is queried against and the secondary site is not started, no results are shown without the report of an SQL error.
- Querying against Multitenant landscapes is limited to single Tenant databases or the system database, meaning there are no views unifying all tenants on the system database similar to the `SYS_DATABASES` schema.

Related Information

[Alert Details \[page 252\]](#)

10.1.2.17 Updating SAP HANA Systems with SAP HANA System Replication

You can use SAP HANA system replication to update your SAP HANA systems.

You can update your SAP HANA systems running in a system replication setup by updating the secondary system first and then updating the primary system. For more information, see *Update an SAP HANA System Running in a System Replication Setup*.

You can use SAP HANA system replication also to upgrade your SAP HANA systems as the secondary system can run with a higher software version than the primary system. For more information, see *Use SAP HANA System Replication for Near Zero Downtime Upgrades*.

Related Information

[Update SAP HANA Systems Running in a System Replication Setup \[page 1203\]](#)

[Use SAP HANA System Replication for Near Zero Downtime Upgrades \[page 1195\]](#)

10.1.2.17.1 Use SAP HANA System Replication for Near Zero Downtime Upgrades

You can use SAP HANA system replication to upgrade your SAP HANA systems as the secondary system can run with a higher software version than the primary system.

Prerequisites

System replication is configured and active between two identical SAP HANA systems:

- The primary system is the production system.
- The secondary system will become the production system after the upgrade.
- The prerequisite is to run both systems with the same endianness.

Context

With system replication active, you can first upgrade the secondary system to a new revision and have it take over in the role of primary system. The takeover is carried out in only a few minutes and committed transactions or data are not lost. You can then do an upgrade on the primary system, which is now in the role of secondary.

i Note

It is possible to reduce the time required to perform an update. For more information, see *Prepare an Update for Flexible System Downtime*.

The secondary system can be initially installed with the new software version or upgraded to the new software version when the replication has already been configured. After the secondary has been upgraded, all data has to be replicated to the secondary system (already having the new software version). When the secondary system is ACTIVE (all services have synced) a takeover has to be executed on the secondary system. This step makes the secondary system the production system running with the new software version.

i Note

If you are upgrading from SAP HANA 1.0 to SAP HANA 2.0, copy the system PKI SSFS key and the data file from the current primary system to the new to-be secondary system. For more information, see SAP Note 2369981 *Required configuration steps for authentication with HANA System Replication*.

In an Active/Active (read enabled) system replication setup the version of the primary and the secondary systems must be identical. For the near zero downtime upgrade to work, the operation mode on the secondary system is automatically set to `logreplay`. Like this the two systems can get back in sync before the takeover step. To re-establish the Active/Active (read enabled) landscape at the end, the operation mode `logreplay_readaccess` must be explicitly specified during the former registration of the primary system as a new secondary system.

For more information about near zero downtime upgrades when using a multitarget system replication setup, see *Use Multitarget System Replication for Near Zero Downtime Upgrades*.

Procedure

1. As `<sid>adm` configure a user in the local userstore under the key SRTAKEOVER. This user requires the necessary privileges to import the repository content of the new version of the software during the takeover process. Use a public host name to access the corresponding SQL port of the System DB (`<SystemDBsqlport>`). Execute this command on the primary and secondary systems:

```
hdbuserstore SET SRTAKEOVER <publichostname>:<SystemDBsqlport> <myrepouser>
<myrepouser_password>
```

i Note

This configuration step should be performed only in the system database, not in every single tenant.

Create a `<myrepouser>` user with the necessary privileges to import the repository content as follows:

```
CREATE USER MY_REPO_IMPORT_USER PASSWORD MyRepoUserPW123;
GRANT EXECUTE ON SYS.REPOSITORY_REST TO MY_REPO_IMPORT_USER;
GRANT REPO.READ ON ".REPO_PACKAGE_ROOT" TO MY_REPO_IMPORT_USER;;
GRANT REPO.IMPORT TO MY_REPO_IMPORT_USER;
GRANT SELECT ON SYS.REPO.DELIVERY_UNITS TO MY_REPO_IMPORT_USER
GRANT REPO.ACTIVATE_IMPORTED_OBJECTS ON ".REPO_PACKAGE_ROOT" TO
MY_REPO_IMPORT_USER
```

For example, for public hostname "mypublichost" and system number "00", "MY_REPO_IMPORT_USER", and "MyRepoUserPW123" :

```
hdbuserstore SET SRTAKEOVER mypublichost:30013 MY_REPO_IMPORT_USER
MyRepoUserPW123
```

The hostname has to be the public host name of the host that the command is executed on and the port is its SQL port number of the SystemDB.

For more information see the section, Secure User Store (hdbuserstore) in the *SAP HANA Security Guide*.

i Note

The command has to be executed on all hosts in a scale-out configuration. If the password for the repository import user is changed the password saved in the userstore also has to be changed.

2. Upgrade the secondary system's SAP HANA server software and all other components.

From your installation directory execute as root:

```
./hdblcm --action=update
```

3. Verify that system replication is active and that all services are in sync.

You can check that the column REPLICATION_STATUS in M_SERVICE_REPLICATION has the value ACTIVE for all services)

4. Stop the primary system.
5. Perform a takeover on the secondary system, including switching virtual IP addresses to the secondary system, and start using it productively.

As `<sid>adm` perform a takeover:

```
hdbnsutil -sr_takeover
```

6. If XS Advanced is being updated as well, update the XS Advanced applications.

```
./hdblcm --action=update
```

7. Upgrade the original primary from the installation directory as root user using the option `--hdbupd_server_nostart` together with all the other components. This is necessary because otherwise the primary has to be stopped again before it can be registered as the secondary.

```
./hdblcm --action=update --hdbupd_server_nostart
```

i Note

For a fast synchronization of the sites – after reregistering the original primary system – perform this failback within the time given by the parameter `datashipping_snapshot_max_retention_time` (default

300 minutes). Otherwise a full data shipping will be done. Furthermore, the optimized resync depends on the availability of the last snapshot. For more information for near zero downtime upgrades in multitier system replication, see *SAP Note 2386973*.

8. Register the original primary as secondary as `<sid>adm`.

```
hdbnsutil -sr_register --name=<secondary_alias>
--remoteHost=<primary_host> --remoteInstance=<primary_systemnr>
--replicationMode=[sync|syncmem|async]
```

9. Start the original primary.

Related Information

[Use Multitarget System Replication for Near Zero Downtime Upgrades \[page 1198\]](#)

[SAP Note 2369981](#)

[SAP Note 1984882](#)

[SAP Note 2386973](#)

[SAP Note 2494079](#)

[SAP Note 2407186](#)

10.1.2.17.1.1 Use Multitarget System Replication for Near Zero Downtime Upgrades

You can upgrade your SAP HANA systems using a multitarget system replication setup.

Prerequisites

System replication is configured and active between identical SAP HANA systems:

- The primary system is the production system.
- The secondary system located in the same data center as the primary system will become the production system after the upgrade.
- There is no replication error.
- The prerequisite is to run all systems with the same endianness.

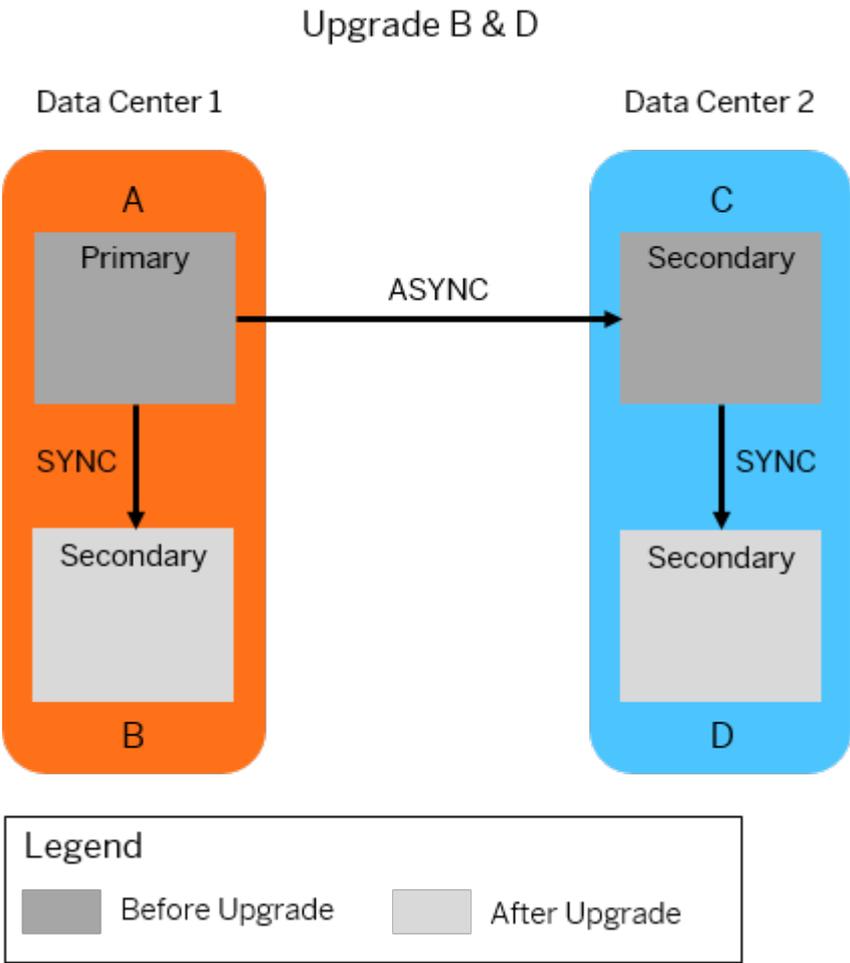
Context

We are using the setup described in *Multitarget System Replication* to exemplify the procedure. In this setup, primary system A replicates data changes to secondary system B located in the same data center. Primary

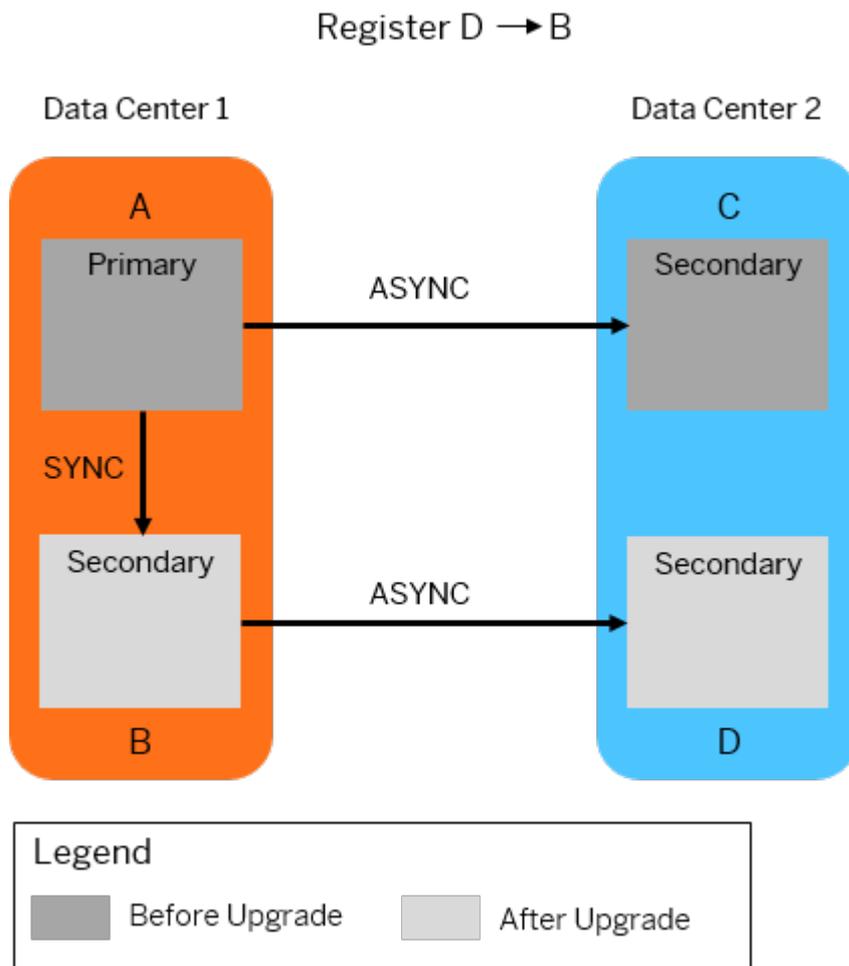
system A also replicates data changes to the secondary system C located in data center 2. Secondary system C is a source system for a further secondary system D located in the same data center with system C.

Procedure

- 1. Upgrade secondary system B in data center 1 and secondary system D in data center 2.

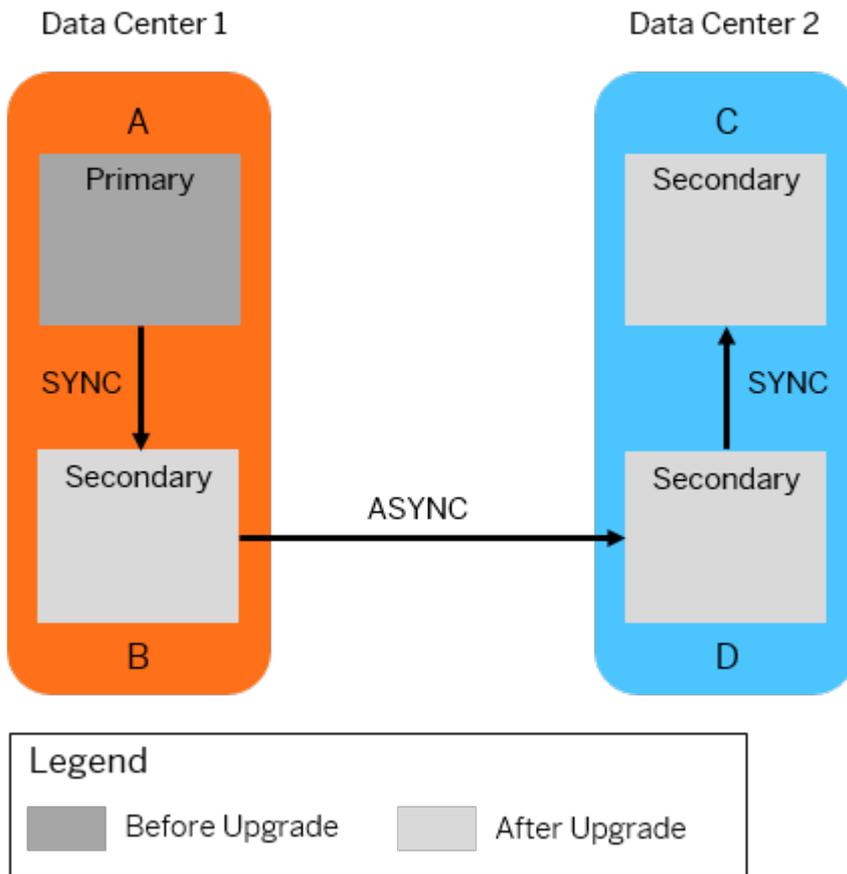


- 2. Register secondary system D in data center 2 to secondary system B in data center 1.



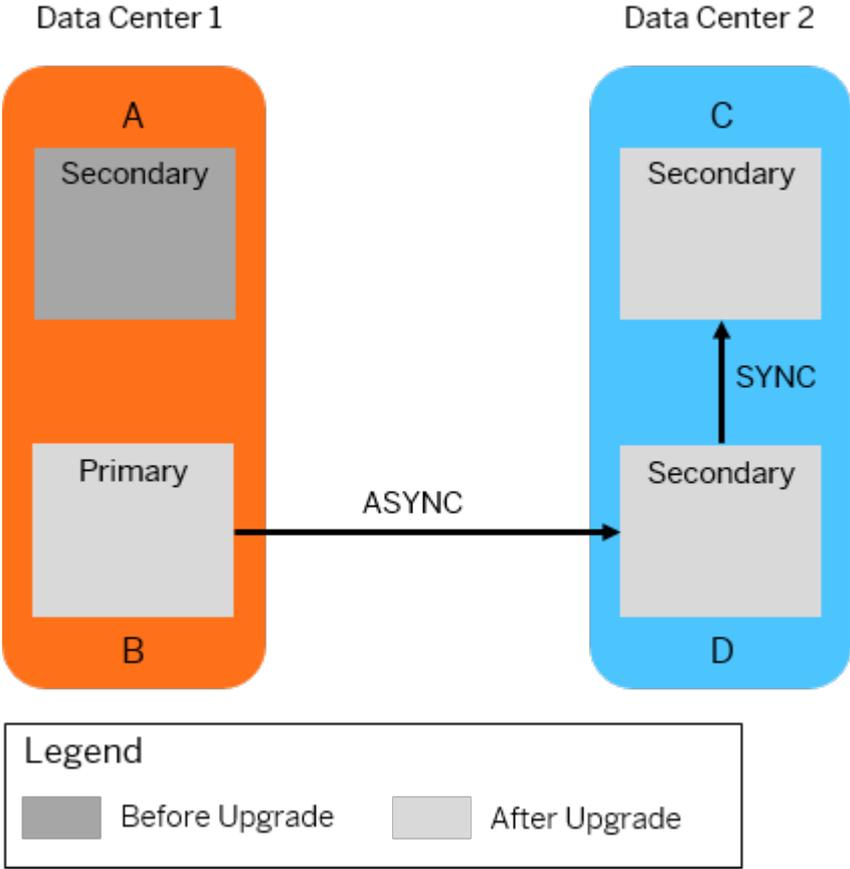
3. Upgrade secondary system C in data center 2. Then, register secondary system C to secondary system D in data center 2.

Upgrade C
Register C → D

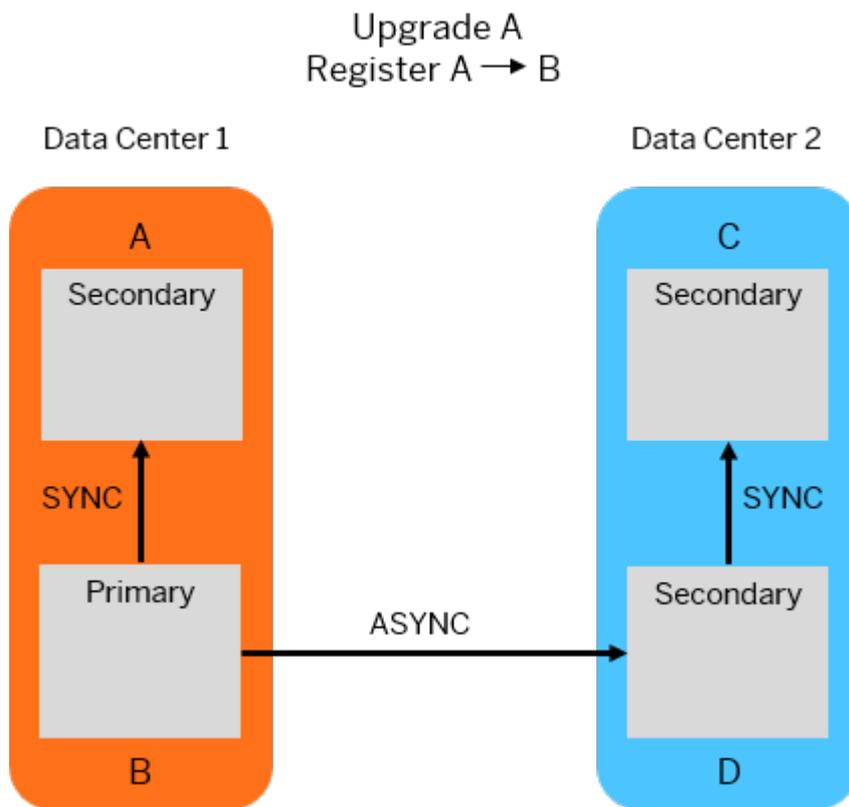


4. Take over on secondary system B in data center 1.
After takeover, secondary system B will be the new primary system.

Takeover on B



- 5. Upgrade and register the previous primary system A to the new primary system B in data center 1.



Related Information

[Multitarget System Replication \[page 1153\]](#)

10.1.2.17.2 Update SAP HANA Systems Running in a System Replication Setup

You can update your SAP HANA system with active system replication by updating the secondary and primary system one after the other.

Prerequisites

- System replication is configured and active between two SAP HANA systems.

Context

You can update your SAP HANA system running in a system replication setup by updating the secondary system first and then updating the primary system.

→ Remember

For system replication setups it is required that at all time the secondary system has the same as or a higher version than the primary system. The secondary system must therefore always be updated before the primary system.

Updating one system after the other results in some downtime. If you want to update your system with reduced downtime, see *Use SAP HANA System Replication for Near Zero Downtime Upgrades*.

i Note

It is possible to reduce the time required to perform an update. For more information, see *Prepare an Update for Flexible System Downtime*.

Procedure

1. Upgrade the SAP HANA server software and all installed components on the secondary system.

From your installation directory execute as root or as `<sid>adm`:

```
./hdblcm --action=update
```

2. With the secondary system online use the SAP HANA lifecycle management tools to upgrade all the other components to the same revision as the server software.
3. Verify that system replication is active and that all services are in sync.

You can check that the column REPLICATION_STATUS in M_SERVICE_REPLICATION has the value ACTIVE for all services)

4. Upgrade the SAP HANA server software and all installed components on the primary system.

From your installation directory execute as root or as `<sid>adm`:

```
./hdblcm --action=update
```

5. With the secondary system online use the SAP HANA lifecycle management tools to upgrade all the other components to the same revision as the server software.
6. Verify that system replication is active and that all services are in sync.

Related Information

[Use SAP HANA System Replication for Near Zero Downtime Upgrades \[page 1195\]](#)

[SAP Note 2407186](#)

10.1.2.18 Add a New Host to a Replicated System

You can add a new host to a replicated system with the SAP HANA lifecycle manager.

Context

i Note

Hosts must be added equally to both primary and secondary sites.

System replication need not be turned off when adding a host.

i Note

It is recommended that a host is added to the secondary site before adding it to the primary site. This avoids the situation where the new host saves data without first being in sync.

Procedure

1. Add a host to the secondary site and start it.
2. Add a host to the primary site and start it.
Replication begins automatically.
3. To remove a host, first remove it from the primary site and then remove the host from the secondary site.

Related Information

[Add Hosts Using the Command-Line Interface \[page 1416\]](#)

[Remove Hosts Using the Command-Line Interface \[page 1424\]](#)

10.1.2.19 Configure Secure Communication (TLS/SSL) Between Primary and Secondary Sites

Configure TLS/SSL on communication channels between primary and secondary systems using the system public key infrastructure (PKI).

Prerequisites

- You have the credentials of the operating system user, <sid>adm.
- You have the system privilege INIFILE ADMIN.

Context

The following communication channels between primary and secondary systems can be secured using TLS/SSL:

- Metadata channel used to transmit metadata (for example, topology information) between the sites
- Data channel used to transmit data between the sites.

i Note

On SAP HANA systems with dynamic tiering, the following steps also enable the system PKI for internal system replication communication. No additional steps are required. Before you configure communication for dynamic tiering see SAP Note 2447994 - *SAP HANA Dynamic Tiering Support for SAP HANA System Replication*.

Be aware that you need additional licenses for SAP HANA options and capabilities. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities*.

Procedure

1. Shut down all systems.

If you want to avoid downtime when enabling TLS/SSL, disable system replication. You can enable or disable TLS/SSL without downtime only if the primary system is not enabled.

2. In the primary and secondary system, enable TLS/SSL for the data channel.

In the `global.ini` file, configure the property `[system_replication_communication] enable_ssl`. The following values are possible:

Value	Description
off (default)	TLS/SSL is disabled for replication source and target systems
on	TLS/SSL is enabled for replication source and target systems

i Note

You must enable SSL for the whole system, that is, in the global.ini file of the system database. Setting this feature for single tenant databases is not supported.

For a simple system replication scenario involving two systems, it is sufficient to set the property to **on** in both systems.

For multitier and multitarget system replication scenarios involving three systems, you can apply **on** in all 3 systems to secure all system replication connections. Alternatively, you can use `site_name` as index to secure either only the communication to the tier 3 secondary system or only the communication to the primary system.

🔗 Example

To exclude the communication between the primary and the secondary, and to secure the communication between all other systems, set the parameter as follows:

```

siteA      ----->      siteB      ----->      siteC
  enable_ssl=on           enable_ssl=on
enable_ssl=on
  enable_ssl[siteB]=off   enable_ssl[siteA]=off

```

i Note

To avoid communication failure between systems, TLS/SSL must be enabled on all systems at the same time. TLS/SSL won't be used unless the secondary system reconnects with the primary. To do this either restart the primary and secondary systems, or re-setup system replication .

3. As `<sid>adm`, restart the `sapstartsrv` service on the secondary system(s):
 - a. `sapcontrol -nr <instance_no> -function StopService`
 - b. `/usr/sap/<sid>/HDB<instance_no>/exe/sapstartsrv pf=/usr/sap/<sid>/SYS/profile/<sid>_HDB<instance_no>_<host> -D -u <sid>adm`
4. Restart all systems.

Related Information

[SAP Note 2447994 - SAP HANA Dynamic Tiering Support for SAP HANA System Replication](#) 

10.1.3 Setting Up Host Auto-Failover

Host auto-failover is a local fault-recovery solution that can be used as a supplemental or alternative measure to system replication. One (or more) standby hosts are added to an SAP HANA system, and configured to work in standby mode.

The databases on the standby hosts do not contain any data and do not accept requests or queries as long as they are in standby mode.

When an active worker host fails, a standby host automatically takes its place. Since the standby host may take over operation from any of the primary hosts, it needs access to all the database volumes. This can be accomplished by a shared networked storage server, by using a distributed file system, or with vendor-specific solutions that can dynamically mount networked storage upon failover.

For more information about how to add an additional standby host, see *Adding Hosts to an SAP HANA System*.

You can monitor the status of all active and standby hosts in the SAP HANA cockpit and the SAP HANA studio.

A HA/DR provider script is available to provide hooks that can be called in response to events during host auto-failover. For more information, see *Implementing a HA/DR Provider*.

Related Information

[Adding Hosts to an SAP HANA System \[page 1413\]](#)

[Configuring Clients for Failover \[page 1208\]](#)

[Configuring Application Servers for Failover \[page 1210\]](#)

[Configure HTTP Load Balancing for SAP HANA Extended Application Services, Classic Model \[page 1211\]](#)

[Implementing a HA/DR Provider \[page 1215\]](#)

[SAP HANA Host Auto-Failover !\[\]\(6f8500c1bdcdf9aa6d3e25cb1d3503b8_img.jpg\)](#)

[SAP HANA Storage Requirements !\[\]\(871c789f3b3190ea47152b4f3a1ca361_img.jpg\)](#)

10.1.3.1 Configuring Clients for Failover

You can configure failover support for clients so that they continue to work in a transparent way to the user in the event of a failover.

SAP HANA clients that were configured to reach the original host need to be sent to the standby host after host auto-failover.

One way to handle this is using a network-based (IP or DNS) approach. Alternatively, SQL/MDX database clients can be configured with the connection information of multiple hosts, optionally including the standby host, by providing a list of hosts in the connection string. The client connection code uses a "round-robin" approach to reconnect, thus ensuring that the client can reach the SAP HANA database, even after failover.

To support failover with client libraries, you have to specify a list of host names separated by a semicolon instead of a single host name. Only hosts that have the role master or standby should be used.

To determine which hosts to use, execute the following SQL statement:

```
SELECT HOST FROM M_LANDSCAPE_HOST_CONFIGURATION WHERE NAMESERVER_CONFIG_ROLE LIKE 'MASTER%' ORDER BY NAMESERVER_CONFIG_ROLE
```

Since one of these master candidates will be active, only they have to be added. When hosts are added to a system, the master list is extended to three hosts, meaning there is one host configured as the actual master and two worker hosts are configured as master candidates. When the first standby host is added to the system, a worker host is removed from this list and replaced by the standby host. This is done because it is faster to fail over to an idle standby host than to an active worker host.

The client will choose one of these hosts to connect to. If a host is not available, the next host from the list will be used. Only in the case that none of the hosts are available will you get a connection error.

If a connection gets lost when a host is not available any longer, the client will reconnect to one of the host specified in the host list.

Example Configurations

Client	Example
JDBC	<pre>Connect URL: jdbc:sap://host1:30015;host2:30015;host3:30015/</pre>
SQLDBC	<pre>SQLDBC_Connection *conn = env.createConnection(); SQLDBC_Retcode rc = conn->connect ("host1:30015;host2:30015;host3:30015", "", "user", "password");</pre>
ODBC	<pre>Connect URL: "DRIVER=HDBODBC32; UID=user; PWD=password; SERVERNODE=host1:30015,host2:30015,host3:30015";</pre>

HTTP Client Access via SAP HANA Extended Application Services, Classic Model

To support HTTP (Web) clients accessing SAP HANA via the SAP HANA XS classic server, it is recommended to install an external, itself fault protected, HTTP load balancer (HLB), such as SAP Web Dispatcher, or a similar product from another vendor. The HLBs are configured to monitor the Web servers on all hosts on all sites. For more information see, *Configuring HTTP Load Balancing for SAP HANA Extended Application Services (XS)*.

If an SAP HANA instance fails, the HLB, which serves as a reverse web-proxy, redirects the HTTP clients to the running SAP HANA XS instance on an active host. HTTP clients are configured to use the IP address of the HLB itself, which is obtained via DNS, and remain unaware of any SAP HANA failover activity.

Related Information

[Client Connection Recovery after Takeover \[page 1132\]](#)

10.1.3.2 Configuring Application Servers for Failover

You can configure failover support for application servers by using the secure user store of the SAP HANA client (`hdbuserstore`) to specify a list of host names that the server can connect to.

For the clients in a host auto-failover landscape, the use of virtual IP addresses is recommended. You can store user logon information, including passwords, in the secure user store of the SAP HANA client (`hdbuserstore`). This allows client programs to connect to the database without having to enter a password explicitly.

The `hdbuserstore` can also be used to configure failover support for application servers (for example, for SAP Business Warehouse) by storing a list of all (virtual) host names to which the application server can connect. All nodes that are master candidates must be added to the `hdbuserstore`.

→ Tip

For more information about how to find out the three master candidates in a distributed system, see SAP Note 1930853.

The application server will choose one of these hosts to connect to from the list. If a host is not available, the next host from the list will be used. Only if none of the hosts are available will you get a connection error. If a connection gets lost when a host is no longer available, the application server will reconnect to one of the hosts specified in the host list.

You can specify a list of host names in the secure user store using the following `hdbuserstore`:

```
hdbuserstore SET default "<hostname_node1>:3<system_number>15; .... <hostname_node(n)>: 3<system_number>15" SAP<sid> <Password>
```

≡ Sample Code

```
hdbuserstore SET default  
"1d9490:33315;1d9491:33315;1d9492:33315;1d9493:33315" SAPP20 <password>
```

KEY default

```
ENV : 1d9490:33315;1d9491:33315;1d9492:33315;1d9493:33315  
USER: SAPP20
```

For more information about `hdbuserstore`, see the *SAP HANA Security Guide*

Related Information

[SAP Note SAP Note 1930853](#)

10.1.3.3 Configure HTTP Load Balancing for SAP HANA Extended Application Services, Classic Model

To enable load balancing for HTTP access to the SAP HANA XS classic sever, you need to set up a load balancer (for example, SAP Web Dispatcher).

Context

To support HTTP (Web) clients accessing SAP HANA via the SAP HANA XS classic server, it is recommended to install an external, itself fault protected, HTTP load balancer (HLB), such as SAP Web Dispatcher, or a similar product from another vendor. The HLBs are configured to monitor the Web servers on all hosts on all sites.

The SAP Web Dispatcher automatically reads the system topology from SAP HANA XS classic and is notified of changes to the topology, for example, when a host is no longer available or a standby host has taken over. The SAP Web Dispatcher then sends requests to a running XS instance on an active host. Third-party load balancers often use a static configuration with an additional server availability check.

The SAP Web Dispatcher can be configured with a list of the three master hosts. Once one of the master hosts is available the SAP Web Dispatcher acquires the topology information. HTTP clients can be configured to use the IP address of the HTTP load balancer itself, and remain unaware of any SAP HANA failover activity.

i Note

For more information about using and configuring the SAP Web Dispatcher for load balancing with SAP HANA multitenant database containers, see *Using SAP Web Dispatcher for Load Balancing with Tenant Databases*.

Procedure

1. Install SAP Web Dispatcher with a minimum release of 7.40 using the SAP NetWeaver Software Delivery Tool and update it to the latest version available on the SAP Software Download Center.
2. Log on to the SAP Web Dispatcher host as the <SID>adm user. Here the <SID> refers to the one of the SAP Web Dispatcher installation.
3. Open the instance profile of your SAP Web Dispatcher.

The SAP Web Dispatcher profile can be found in the following location:

```
usr/sap/<SID>/SYS/profile
```

4. Disable the ABAP system configuration, which is done automatically during the installation by commenting out the entries in this section of the profile:

```
# Accessibility of Message Servers
-----
#rdisp/mshost = ldcialx
#ms/http_port = 8110
```

5. Add a list of semicolon separated URLs and the base URL (without path) used for fetching topology information, to the XSSRV parameter in the profile.

An example could be:

```
wdisp/system_0 = SID=HDX, XSSRV=http://1d9490:8089;http://1d9491:8089,  
SRCSRV=*
```

Related Information

[SAP Web Dispatcher](#)

[SAP Note 1855097](#)

[SAP Note SAP Note 1883147](#)

[Using SAP Web Dispatcher for Load Balancing with Tenant Databases \[page 280\]](#)

10.1.3.4 Host Auto-Failover Parameters

This topic provides an overview of parameters available for Host Auto-Failover.

Master Failover without Standby Hosts

Distributed landscapes without standby hosts may also perform a failover to ensure that the master host is always available.

Parameter	<code>nameserver.ini/[failover]/enable_master_failover</code>
------------------	---

Description:	When set to false, the masterize check of the nameserver master candidates is disabled. Furthermore, adding a new host does not modify the master candidates list.
--------------	--

Online	Yes
--------	-----

Change:

Default:	True
----------	------

Failover Groups

During installation a failover group can be configured per host. If a failover target host is available in the same group, it will be preferred over hosts from other groups. This can be used to achieve better "locality" in large systems (for example, to use network or storage connection with less latency). It can also be used to separate differently sized hardware or storages.

Parameter `nameserver.ini/[failover]/cross_failover_group`

Description: When set to false, failover is restricted to the hosts in the same group.

Online Yes

Change:

Default: False

Automatic Host Shutdown by Service Failures

For every service a fixed number of restarts can be defined after which the daemon stops. The nameserver is the only service that has the first parameter set to true as default. This means that any problem involving a constant nameserver crash will eventually stop the daemon.

Parameter `daemon.ini/[failover]/startup_error_shutdown_instance`

Description: When set to true, the daemon will shut down all services on the host if this service cannot start.

Online Yes

Change:

Default: True

Parameter `daemon.ini/[failover]/startup_error_restart_retries`

Description: The number of retries if a service fails in the startup procedure.

Online Yes

Change:

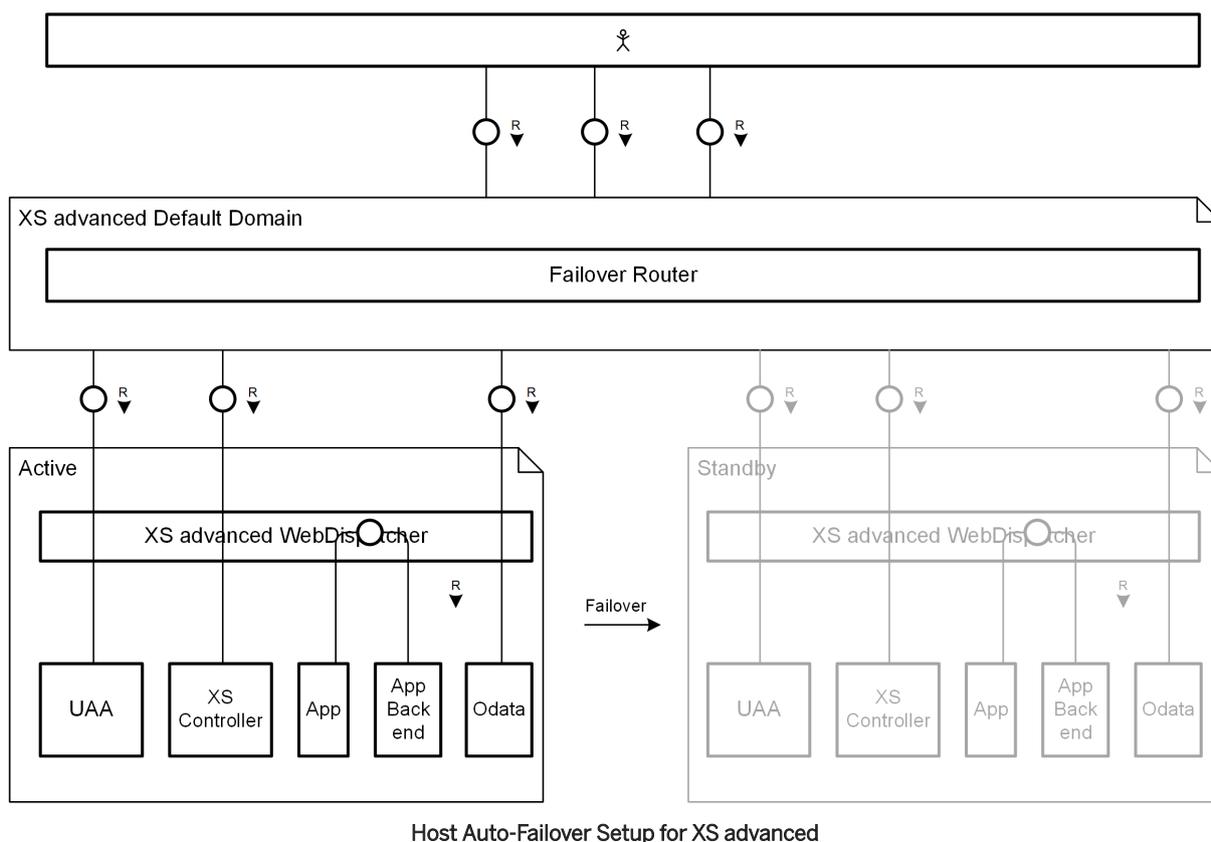
Default: 4

10.1.3.5 Host Auto-Failover Setup with XS Advanced Runtime

This topic describes the host auto-failover setup for XS advanced runtime.

The host auto-failover setup for XS advanced is similar to the setup of XS advanced behind a reverse proxy. For more information on the XS advanced setup behind a reverse proxy, see *Setting Up the XS Advanced Runtime Behind a Reverse Proxy* in the *SAP HANA Server Installation and Update Guide*.

In the host auto-failover setup of XS advanced the reverse proxy component (referred to as failover router) is routing to either the active system or the standby system after takeover.



Host Auto-Failover Setup for XS advanced

Hosts with a standby role can be added to the SAP HANA system. These hosts will take over in case the active host fails.

New host roles are introduced with XS advanced runtime: `xs_worker` and `xs_standby`.

If a host already has a worker role and an `xs_worker` role, it's sufficient to configure a failover host with a standby role, which will take over both these roles in case the host fails.

In case a host only has an `xs_worker` role, but no worker role, you must configure a host with the role `xs_standby` for the takeover.

Certificate Setup

A single certificate is used for both the active and the standby XS advanced Webdispatcher. If you terminate SSL at the reverse proxy, the failover router must trust both the active and the standby hosts. This can be established in two ways:

- Create a certificate that contains all potential master hosts (active and standby) in the *Subject Alt Names* section. Then, install it on the XS advanced WebDispatcher.
Or
- Turn off the hostname validation in the failover router after checking if this has any security implications.

For more information on setting up certificates, see the sections *Certificate Setup without SSL Termination at Reverse Proxy* and *Certificate Setup with SSL Termination at Reverse Proxy* in *Establishing the Trust Relationship with a Reverse Proxy*.

10.1.4 Implementing a HA/DR Provider

The SAP HANA nameserver provides a Python-based API, which is called at important points of the host auto-failover and system replication takeover process.

These so called "hooks" or HA/DR providers can be used for arbitrary operations that need to be executed. One of the most important uses of the failover hooks is moving around a virtual IP address (in conjunction with STONITH).

Nevertheless, there are other purposes like starting tools and applications on certain hosts after failover or even stopping DEV or QA SAP HANA instances on secondary sites before takeover. Multiple failover hooks can be installed and used in parallel with a defined execution order.

i Note

When calling subprocesses within a HA/DR Provider implementation, please refrain from using the Python modules `subprocess` and `popen2`, as well as `os.popen2`, `os.popen3`, and `os.popen4`. Those methods allocate memory, which can cause a deadlock when forking the nameserver process. Use `os.popen()` or `os.system()` instead, even though it is marked as deprecated.

This section will describe the Python API in detail and gives an example how the hooks can be used. An implemented Python class is called a HA/DR provider. This script contains hook methods, which are called at certain events.

The failover hooks are included in SAP HANA. SAP HANA comes with its own Python interpreter, which is used for interpreting the user defined failover hooks. The failover hook API also has a version number.

10.1.4.1 Create a HA/DR Provider

You can adapt Python files delivered with SAP HANA to create your own HA/DR provider. This allows you to integrate, for example, SAP HANA failover mechanisms into your existing scripts.

Context

To create your own HA/DR provider, the following steps must be executed and then add the methods you want to use from those listed in *Hook Methods*:

Procedure

1. Create a new directory for the HA/DR provider

The directory should be within the shared storage of the SAP HANA installation, but outside the <SID> directory structure (otherwise it is likely to be deleted/overwritten during a SAP HANA update).

For example you could use the following location: `/hana/shared/myHooks`

- Copy the `exe/python_support/hdb_ha_dr/HADRDummy.py` from an installed SAP HANA system to the new location.
For example, copy the file to: `/hana/shared/myHooks/myFirstHook.py`

Note

Do not copy the `client.py`, otherwise updates and new features will be missed when updating SAP HANA. When using the import statement as described below the `client.py` from the SAP HANA installation will be used.

- Adapt the contents of the new file by renaming the Python class to the name of the file, for an example see Code Listing 1
- Fill out the Python dict in the `about()` method, for an example see Code Listing 1.

Sample Code

Code Listing 1

```
from hdb_ha_dr.client import HADRBase, Helper
import os, time
class myFirstHook(HADRBase):
    apiVersion = 1
    def __init__(self, *args, **kwargs):
        # delegate construction to base class
        super(myFirstHook, self).__init__(*args, **kwargs)
    def about(self):
        return {"provider_company" :      "SAP",
                "provider_description" :  "Template Dummy Provider",
                "provider_version" :      "1.0"}
```

Within the SAP HANA environment, the path `exe/python_support` is part of the `PYTHONPATH` setting. Therefore, `hdb_ha_dr` can be used as module for the import of the base class in helper class (shown in the first line of Code Listing 1).

The attribute `apiVersion` defines which HA/DR provider API version will be used.

The `__init__()` method should always call the super method with the parameters `*args, **kwargs` in order to ensure the correct initialization of the tracer and configuration file wrapper. If required, additional initialization steps can be used here.

Finally, the `about()` method must return a Python dict with the keys as shown. The values will be used for monitoring in the `M_HA_DR_PROVIDERS` view.

There are three class attributes defined in the HA/DR provider base class::

- `tracer`: a tracer that is available to all derived classes tracing to the nameserver's trace file
- `config`: a wrapper for easy access of optional configuration parameters in the `global.ini`
- `apiVersion`: the definition of the API version

More details can be found in the sections `Additional Configuration Parameters` and `Tracing` (see `Related Information`).

Next Steps

With the basic HA/DR provider now implemented you can continue by choosing and adding the methods listed in *Hook Methods* to your provider.

Related Information

[Hook Methods \[page 1217\]](#)

[Install and Configure a HA/DR Provider Script \[page 1223\]](#)

10.1.4.1.1 Hook Methods

There are a number of pre takeover, post takeover and general hooks available for you to use.

The following hook methods are available:

Name	Trigger	Caller	Landscape	Error behavior
startup()	Beginning of name-server's start up phase	Starting nameserver	Each individual host	Nameserver aborts, start up is canceled
shutdown()	Just before the name-server exists	Stopping nameserver	Each individual host	Error trace is written
failover() [Host Auto-Failover]	As soon as the name-server made a decision about the new role	Host that takes over the role	Each host that gets a new role	Nameserver aborts, failover is canceled
stonith() [Host Auto-Failover]	As soon as the name-server made the decision about the new role	Master nameserver	For each failed host	Nameserver aborts, failover is canceled
preTakeover() [System Replication]	As soon as the hdbnsutil -sr_takeover command is issued	Master nameserver	Called only once on the master	Takeover is aborted
postTakeover() [System Replication]	As soon as all services with a volume return from their assign-call (open SQL port)	Master nameserver	Called only once on the master	Error trace is written
srConnection-Changed()	As soon as one of the replicating services loses or (re-)establishes the system replication connection	Master nameserver	Called only once on the master	Error trace is written

Name	Trigger	Caller	Landscape	Error behavior
srServiceState-Changed() [System Replication]	As soon as the nameserver made a decision about the new state	Host that detects a local service change	Each individual host	Error trace is written
srReadAccessInitialized() [System Replication]	As soon as a tenant database or the system database is ready to accept SQL read queries on a read enabled secondary system	Master node on the primary system	Called only once on the master node on the primary system	Error trace is written

All hook methods receive a set of parameters, which can be used to identify the state and configuration of the calling nameserver. The nameserver expects the return code 0 in case of successful execution and codes other than 0 for the error case.

The hook methods are shown in detail in the following sections.

Startup hook method stub

Code Listing 3 shows the startup hook method stub.

```
def startup(self, hostname, storage_partition, system_replication_mode,
**kwargs):
    """
    Hook description:

    * time of call: beginning of startup of the nameserver
    * caller: the starting host
    * landscape: each host calls it individually
    * behavior upon failure: nameserver aborts, startup is cancelled
    @param hostname: the local hostname
    @type hostname: string
    @param storage_partition: the storage partition number, 0 for standby
hosts
    @type storage_partition: int
    @param system_replication_mode: mode of system replication
    @type system_replication_mode: string
    @param **kwargs: place holder for later usage (new parameters) to
        keep the interface stable
    @type **kwargs: dict
    @return: information about success
    @rtype: int
    """
    return 0
```

Shutdown hook method stub

Code Listing 4 shows the shutdown hook method stub.

i Note

The invocation of the `shutdown()` method is not guaranteed. If the nameserver is terminated prematurely it cannot call the HA/DR provider, for example if a host fails or during SAP HANA shutdown, when the services have not had enough time to run through their shutdown routines. Therefore crucial tasks such as the deletion of virtual IPs need to be implemented in the `startup()` and `failover()` method.

```
def shutdown(self, hostname, storage_partition, system_replication_mode,
**kwargs):
    """
    Hook description:

    * time of call: just before the nameserver exists
    * caller: the stopping host
    * landscape: each host calls it individually
    * behavior upon failure: error trace is written
    @param hostname: the local hostname
    @type hostname: string
    @param storage_partition: the storage partition number, 0 for standby
hosts
    @type storage_partition: int
    @param system_replication_mode: mode of system replication
    @type system_replication_mode: string
    @param **kwargs: place holder for later usage (new parameters) to
        keep the interface stable
    @type **kwargs: dict
    @return: information about success
    @rtype: int
    """
    return 0
```

Failover hook method stub

Code Listing 5 shows the failover hook method stub.

```
def failover(self, hostname, storage_partition, system_replication_mode,
**kwargs):
    """
    Hook description:

    * time of call: when the nameserver made the decision about the new role
    * caller: the host that takes over the role
    * landscape: called on each host that gets a new role
    * behavior upon failure: nameserver aborts, failover is cancelled
    @param hostname: the local hostname
    @type hostname: string
    @param storage_partition: the storage partition number, 0 for standby
hosts
    @type storage_partition: int
    @param system_replication_mode: mode of system replication
    @type system_replication_mode: string
    @param **kwargs: place holder for later usage (new parameters) to
        keep the interface stable
    @type **kwargs: dict
    @return: information about success
    @rtype: int
    """
    return 0
```

stonith hook method stub

Code Listing 6 shows the stonith hook method stub.

```
def stonith(self, failing_host, **kwargs):
    """
    Hook description:

    * time of call: when the nameserver made the decision about the new role
    * caller: the master host
    * landscape: for each failed host
    * behavior upon failure: nameserver aborts, failover is cancelled
    @param failing_host: the SAP HANA internal name of the failed host
    @type failing_host: string
    @param **kwargs: place holder for later usage (new parameters) to
                    keep the interface stable
    @type **kwargs: dict
    @return: information about success
    @rtype: int
    """
    return 0
```

preTakeover hook method stub

Code Listing 7 shows the preTakeover hook method stub.

```
def preTakeover(self, isForce, **kwargs):
    """
    Hook description:

    * time of call: as soon as the hdbnsutil -sr_takeover command is issued
    * caller: the master host
    * landscape: called only once on the master
    * behavior upon failure: nameserver aborts, takeover is cancelled
    @param isForce: flag if it is a normal or forced takeover (as of today,
                    takeover is always forced regardless of the value of the
                    flag)
    @type isForce: bool
    @param **kwargs: place holder for later usage (new parameters) to
                    keep the interface stable
    @type **kwargs: dict
    @return: information about success
    @rtype: int
    """
    return 0
```

postTakeover hook method stub

Code Listing 8 shows the postTakeover hook method stub.

```
def postTakeover(self, rc, **kwargs):
    """
    Hook description:

    * time of call: as soon as all services with a volume return from their
```

```

        assign-call (open SQL port)
    * caller: the master host
    * landscape: called only once on the master
    * behavior upon failure: error trace is written
    @param rc: the return code of the actual takeover process; 0=success,
        1=waiting for forced takeover, 2=failure
    @type rc: int
    @param **kwargs: place holder for later usage (new parameters) to
        keep the interface stable
    @type **kwargs: dict
    @return: information about success
    @rtype: int
    """
    return 0

```

srConnectionChanged hook method stub

Code Listing 9 shows the srConnectionChanged hook method stub.

```

def srConnectionChanged(self, parameters, **kwargs):
    """
    Hook description:
    * time of call: as soon as one of the replicating services loses or
        (re-)establishes the system replication connection
    * caller: master node on primary site
    * landscape: called only once on the master node on primary site
    * behavior upon failure: error trace is written
    * Possible return codes:
    * 0: Ok - continue processing
    * 1: Block - Further SAP HANA processing is blocked. Every 5 sec. there
        will be a retry to call this hook
    * If an HA/DR Provider shall not block SAP HANA processing or if a
    blocking
        situation caused by an HA/DR Provider shall (temporarily) be
        resolved
    * use 'hdbnsutil -sr_blockonconnectionchanged --disable'
    @param parameters: dict of parameters {hostname:string, port:string,
        database:string, status:int, database_status:int,
        system_status:int, timestamp:string, is_in_sync:bool,
    reason:string}
    @type parameters: dict
    @param **kwargs: place holder for later usage (new parameters) to keep
    the
        interface stable
    @type **kwargs: dict
    @return: information about success (0 = continue transaction, 1 = halt
        further transactions)
    @rtype: int
    * parameters:
    * -- hostname: host where the service is running
    * -- port: service's port
    * -- database: service's tenant database (MDC)
    * -- status: service replication status (10: NoHSR, 11: Error, 12:
    Unknown, 13: Initializing, 14: Syncing, 15: Active)
    * -- database_status: tenant database replicating status (10: NoHSR, 11:
    Error, 12: Unknown, 13: Initializing, 14: Syncing, 15: Active)
    * -- system_status: HANA database overall replicating status (10: NoHSR,
    11: Error, 12: Unknown, 13: Initializing, 14: Syncing, 15: Active)
    * -- timestamp: date and time of the event
    * -- is_in_sync: true if service is in sync
    * -- reason: additional details (e.g. 'Starting', 'Stopping')
    * -- siteName: name of the replicating secondary site
    """

```

```
"""
return 0
```

srServiceStateChanged hook method stub

Code Listing 10 shows the srServiceStateChanged hook method stub.

```
def srServiceStateChanged(self, parameters, **kwargs):
    """
    Hook description:
    * time of call: as soon as the nameserver made a decision about the new state
    * caller: host that detects a local service change
    * landscape: each individual host
    * behavior upon failure: error trace is written
    @param parameters: dict of parameters {hostname:string, service_name:string,
    service_port:string, service_status:string, timestamp:string}
    @type parameters: dict
    * parameters:
    * -- hostname: host where the service state has changed
    * -- service_name: name of the service
    * -- service_port: port of the service
    * -- service_status: (no, yes, unknown, starting, stopping)
    * -- timestamp: date and time of the service change event
    """
    return 0
```

srReadAccessInitialized hook method stub

Code Listing 11 shows the srReadAccessInitialized hook method stub.

```
def srReadAccessInitialized(self, parameters, **kwargs):
    """
    Hook description:
    * time of call: when a tenant database or the SystemDB is ready to
    accept SQL read queries on a read enabled secondary system
    * caller: master node on primary site
    * landscape: called only once on the master node primary site
    * behavior upon failure: error trace is written
    @param parameters: dict of parameters {last_initialized_database:string,
    databases_with_read_access_initialized:list<string>,
    databases_without_read_access_initialized:list<string>, timestamp:string,
    all_databases_initialized:bool}
    @type parameters: dict
    @param **kwargs: place holder for later usage (new parameters) to keep
    the interface stable
    @type **kwargs: dict
    * parameters:
    * -- last_initialized_database: tenant database that has finished its
    read access initialization and is ready to accept SQL requests
    * -- databases_with_read_access_initialized: list of tenant databases
    that have finished their read access initialization
    * -- databases_without_read_access_initialized: list of tenant databases
    that have NOT finished their read access initialization
    * -- timestamp: date and time of the event
    * -- all_databases_initialized: true if all tenant databases have
    finished the read access initialization
```

```
""  
return 0
```

10.1.4.2 Install and Configure a HA/DR Provider Script

You can add, configure, and monitor your custom provider scripts in the SAP HANA studio.

If the HA/DR provider script is created, it can easily be installed on a SAP HANA system by adding a section called [ha_dr_provider_<classname>] to the global.ini with following parameters:

- provider : the class name
- path : location of the script
- execution_order : the ordering of the HA/DR provider if there is more than one; this is a number between 1 and 99

An example is shown in Code Listing 9.

Sample Code

Code Listing 9

```
[ha_dr_provider_myfirsthook]  
provider = myFirstHook  
path = /hana/shared/myHooks  
execution_order = 50
```

It is possible to specify multiple HA/DR providers by adding multiple sections.

All scripts are loaded during the start up phase of the name server.

Additional Configuration Parameters

If the HA/DR provider requires additional configurations parameters, arbitrary key value pairs can be added to the configuration parameter section. An example is shown in Code Listing 10.

Sample Code

Code Listing 10 - HA/DR Provider section and custom configuration parameters

```
[ha_dr_provider_myfirsthook]  
provider = myFirstHook  
path = /hana/shared/myHooks  
execution_order = 50  
myparameter1 = somevalue  
myparameter2 = 42
```

To consume these parameters, the configuration parameter wrapper HADRBBaseConfiguration (initialized in base class of the HA/DR Provider) can be used with following methods:

- self.config.hasKey(<name>)

- `self.config.get(<name>)`

Sample Code

Code Listing 11 - Using the configuration parameter wrapper

```
def startup(self, hostname, storage_partition, system_replication_mode,
**kwargs):
    if self.config.hasKey("myparameter1"):
        self.tracer.debug("param2 is '%s'" %
self.config.get("myparameter2"))
    return 0
```

Access Rights

In many cases, the HA/DR provider needs additional rights granted in order to run operating system command and programs that require root access. Usually those rights are granted by adding a line to the `/etc/sudoers` file similar to this:

```
<sid>adm ALL= NOPASSWD: /path/command, /path2/command2
```

For example: `mmtadm ALL= NOPASSWD: /sbin/arping, /sbin/ip`

Execution Order

The order of execution is defined with the `execution_order` parameter in the `ha_dr_<classname>` section of the `global.ini` file by specifying a number between 1 and 99 – the lower the number, the higher the priority. For example:

Sample Code

Code Listing 12 - Configuration for two HA/DR Providers

```
[ha_dr_provider_mySTONITH]
provider = mySTONITH
path = /hana/shared/myHooks
execution_order = 50
[ha_dr_provider_vIPMover]
provider = vIPMover
path = /hana/shared/myHooks
execution_order = 51
```

With the `execution_order` parameter, you can ensure that the `mySTONITH` provider is always called before the `vIPMover` provider

Monitoring with M_HA_DR_PROVIDERS

The monitoring view `M_HA_DR_PROVIDERS` contains all information about installed HA/DR Providers.

	PROVIDER_NAME	PROVIDER_COMPANY	PROVIDER_DESCRIPTION	PROVIDER_VERSION	PROVIDER_TYPE	PROVIDER_PATH	EXECUTION_ORDER
1	vIPMover	SAP	vIP Mover	1.0	GENERIC	/hana/shared/myHooks	51

Tracing

There are a number of methods provided with the HA/DR provider base class that allow you to implement trace levels:

- `self.tracer.debug(<text>)`
- `self.tracer.info(<text>)`
- `self.tracer.warning(<text>)`
- `self.tracer.error(<text>)`
- `self.tracer.fatal(<text>)`

Everything will be traced to the component `ha_dr_<classname>`, in this example it would be `ha_dr_myfirsthook`. The default trace level is "info". You can override the level used by setting the parameter `ha_dr_<classname>=<level>` in the trace section of the `global.ini` file.

Additionally, the name server itself traces general information about the HA/DR provider calls and return code to the trace component `ha_dr_provider`. The default trace level is "info" as well.

10.1.4.3 Example HA/DR Provider Implementation

A full example showing how two HA/DR providers can be implemented for a sample landscape consisting of two SAP HANA systems with system replication enabled.

i Note

This example does not make any statement or assumption about what kind of hardware and software set up is licensed and if it fulfills production SAP HANA requirements at all. This is only a showcase, based on virtual machines, which was available during the development of this example and should not be considered for production use. Concepts such as virtual IPs are not always applicable since the network architecture within and across data centers needs to be considered.

The usage of virtual IPs to automatically reconnect to the master host after a failover or a system replication takeover only works if the virtual IP on the failing host is disabled through a controlled shut down. In a split brain situation (network problems separate parts of the landscape) or on host failures it cannot be guaranteed that the virtual IP is unique in the network causing severe routing issues. Therefore, for some use cases virtual hostnames or cluster manager software to control the assignment of virtual IPs might be an alternative.

The context of this example is based on two SAP HANA systems consisting of 16 virtual machines each. Two HA/DR providers are going to be implemented:

- A provider that sets up a virtual IP address on the master host of the primary system allowing all clients to use always the same IP address for connecting regardless of any HA or DR activities. This HA/DR provider will be called *vIPMover*.
- A provider that runs STONITH in order to ensure proper I/O fencing. STONITH will be called for host auto-failover (the failed host) and system replication takeover (all three master candidates of the other site). The latter one is usually the task of an external cluster manager, but for this simple example, we use the direct way (which is usually not possible in real data center set ups). This HA/DR provider will be called *mySTONITH*.

vIPMover HA/DR Provider

The purpose of this provider is to set up a virtual IP address every time the active master host moves. This can either happen when host auto-failover occurs or by a system replication takeover. For the failover case, the IP address move will simplify the SQL client connect by just having one IP address/hostname to specify. And for the system replication case, the client is able to find seamlessly the new system. However, proper fencing is a crucial part of moving an IP address around in order to avoid split-brain situations, because two hosts listen on the same address. The solution for this problem in this example will be the *mySTONITH* HA/DR provider, which will reboot the virtual machine, which has failed.

The HA/DR provider will make use of the Linux operating system commands `/sbin/arping` and `/sbin/ip`, which need to be added to the `/etc/sudoers` file for SAP HANA's `<sid>adm` user.

Listing 13 shows the class definition, constructor and the `about()` method of the *vIPMover* class.

Sample Code

Code Listing 13 - Class definition, constructor and `about()` method of the *vIPMover* class

```
from hdb_ha_dr.client import HADRBase
import os
class vIPMover(HADRBase):
    apiVersion = 1
    def __init__(self, *args, **kwargs):
        super(vIPMover, self).__init__(*args, **kwargs)
        self.vIP = self.config.get("vip")
        self.eth = self.config.get("eth")
        self.netMask = self.config.get("netmask")
    def about(self):
        return {"provider_company" :      "SAP Documentation Example",
                "provider_description" :  "vIP Mover",
                "provider_version" :     "1.0"}
```

Using `apiVersion = 1`, the `__init__()` method calls its super method and additionally reads the three attributes `vIP`, `eth` and `netMask` from the configuration file. More details later. The next step is to define helper methods for the actual setup of the virtual IP address. This example uses the standard Linux command `ip` and `arping` to set up and shut down an IP address:

Sample Code

Helper methods for virtual IP address setup and shut down

```
def setupIP(self):
    # setup IP
```

```

        command1 = "sudo /sbin/ip addr add %s/%s dev %s" % (self.vIP,
self.netMask, self.eth)
        rc1 = os.system(command1)
        self.tracer.info("command '%s' returned with rc=%s" % (command1, rc1))
        command2 = "sudo /sbin/arping -U -c 5 %s" % self.vIP
        rc2 = os.system(command2)
        self.tracer.info("command '%s' returned with rc=%s" % (command2, rc2))
        return rc1 + rc2
    def shutdownIP(self):
        command = "sudo /sbin/ip addr del %s/%s dev %s" % (self.vIP,
self.netMask, self.eth)
        rc = os.system(command)
        self.tracer.info("comand '%s' returned with rc=%s" % (command, rc))
        return rc

```

The commands that will be executed in this example would be:

```

sudo /sbin/ip addr add 10.208.155.179/20 dev eth0
sudo /sbin/arping -U -c 5 10.208.155.179

```

Finally, Listing 15 shows the implementation of the hook methods.

Sample Code

Code Listing 15 - The hook method implementation

```

    def startup(self, hostname, storage_partition, system_replication_mode,
**kwargs):
        self.shutdownIP()
        # only setup vIP on the primary system
        if system_replication_mode not in ["", "primary"]:
            return 0
        # only setup vIP on the master host
        if storage_partition == 1:
            return self.setupIP()

        return 0
    def shutdown(self, hostname, storage_partition, system_replication_mode,
**kwargs):
        if system_replication_mode not in ["", "primary"]:
            return 0
        if storage_partition == 1:
            return self.shutdownIP()
        return 0
    def failover(self, hostname, storage_partition, system_replication_mode,
**kwargs):
        if system_replication_mode not in ["", "primary"]:
            return 0
        if storage_partition == 1:
            return self.setupIP()
        return 0
    def preTakeover(self, isForce):
        """Pre takeover hook."""
        return self.setupIP()

```

The three methods `startup()`, `shutdown()` and `failover()` have the same structure. The condition `if system_replication_mode not in ["", "primary"]` checks if there is no system replication configured or if it is a system replication primary system.

In the `preTakeover()` method, the virtual IP address is started just before the internal takeover process begins.

The HA/DR provider is configured in the `global.ini` with the three user-defined parameters:

Sample Code

Code Listing 16 - The vIPMover configuration in the SAP HANA system

```
[ha_dr_provider_vIPMover]
provider = vIPMover
path = /hana/shared/myHooks
execution_order = 51
vip = 10.208.155.179
eth = eth0
netmask = 20
```

mySTONITH HA/DR Provider

The second HA/DR provider offers STONITH to the SAP HANA system. For this example the virtual machine sends a hard reboot command with a locally installed API to a management node outside the SAP HANA system. As the system is installed with virtual hostnames, but the API requires the public hostnames, a mapping of both is defined in the global.ini. Another option would be to resolve those names via naming convention if applicable.

Listing 17 shows the HA/DR Provider implementation.

Sample Code

Code Listing 17 - Implementation of the mySTONITH HA/DR provider

```
from hdb_ha_dr.client import HADRBase
import os
class mySTONITH(HADRBase):
    apiVersion = 1
    def __init__(self, *args, **kwargs):
        super(mySTONITH, self).__init__(*args, **kwargs)
        self.hsrStonith = self.config.get("hsr_stonith").split()
    def about(self):
        return {"provider_company" : "SAP Documentation Example",
                "provider_description" : "Basic virtual machine STONITH",
                "provider_version" : "1.0"}

    def stonith(self, failingHost, **kwargs):
        vmName = self.config.get("map_%s" % failingHost)
        if vmName == "":
            raise Exception("hostname for virtual machine not configured")
        self.tracer.info("calling STONITH for %s (%s)" % (vmName,
failingHost))
        return os.system(("<script> <logon credentials> --name %s <hard
reboot> <servicing host>" % (vmName)))
    def preTakeover(self, isForce):
        rc = 0
        for h in self.hsrStonith:
            rc = rc + self.stonith(h)
        return rc
```

The `__init__()` and `about()` methods are filled similar to the example above. The `stonith()` method looks up the SAP HANA internal hostname via configuration parameter and executes the STONITH command. This is specific to the type of the virtual environment. For bare metal servers, an IPMI-based call is a typical implementation.

The `preTakeover()` method sends a STONITH command to all master hosts of the other site defined via configuration parameter `hsr_stonith`, a space-separated list of host names.

As a result the configuration entries look like this:

Sample Code

```
[ha_dr_provider_mySTONITH]
provider = mySTONITH
path = /hana/shared/myHooks
execution_order = 50
hsr_stonith = hananode17 hananode28 hananode32
map_hananode17 = DEWDFTVU3017
map_hananode28 = DEWDFTVU3028
map_hananode32 = DEWDFTVU3032
map_hananode01 = DEWDFTVU3001
map_hananode02 = DEWDFTVU3002
map_hananode03 = DEWDFTVU3003
map_hananode04 = DEWDFTVU3004
map_hananode05 = DEWDFTVU3005
map_hananode06 = DEWDFTVU3006
map_hananode07 = DEWDFTVU3007
map_hananode08 = DEWDFTVU3008
map_hananode09 = DEWDFTVU3009
map_hananode10 = DEWDFTVU3010
map_hananode11 = DEWDFTVU3011
map_hananode12 = DEWDFTVU3012
map_hananode13 = DEWDFTVU3013
map_hananode14 = DEWDFTVU3014
map_hananode15 = DEWDFTVU3015
map_hananode16 = DEWDFTVU3016
```

10.2 SAP HANA Database Backup and Recovery

SAP HANA offers comprehensive functionality to safeguard your database and ensure that it can be recovered speedily and with maximum business continuity.

i Note

This documentation only covers backup and recovery of an SAP HANA database. It does not describe how to back up and recover all the components that can be part of an SAP system.

SAP HANA supports the following backup and recovery capabilities:

- Full backups
 - Data backups
 - Data snapshots

i Note

To make use of storage snapshot-based SAP HANA backups, first create an SAP HANA data snapshot.

- Delta backups

- Incremental backups
- Differential backups
- Redo log backups
- Backup and recovery using third-party tools
- Integrity checks for backups
- Backup lifecycle management
- Recovery to a point-in-time
- Recovery to a specific data backup or data snapshot (without using the log area or log backups)
- Database copy using backup and recovery

Related Information

[SAP HANA Backup Types \[page 1246\]](#)

[Working with Third-Party Backup Tools \[page 1303\]](#)

[Manually Checking Whether a Recovery is Possible \[page 1335\]](#)

[Recovering an SAP HANA Database \[page 1347\]](#)

[Copying a Database Using Backup and Recovery \[page 1374\]](#)

10.2.1 Savepoints and Redo Logs

To maintain optimal performance, an SAP HANA database holds the bulk of its data in memory. However, SAP HANA also uses persistent storage to provide a fallback in the event of a fault or a failure.

During normal database operation, changed data is automatically saved from memory to disk at regular **savepoints**. By default, savepoints are created every five minutes, including during a backup.

With a system running on properly configured hardware, the impact on performance of savepoints is negligible. Savepoints do not affect the processing of transactions. During a savepoint, transactions continue to run as normal, and new transactions can be started as normal.

Additionally, all data changes are recorded in the **redo log buffer**. When a database transaction is committed, the redo log buffer is saved to disk. Also, if the redo log buffer fills at any time, the redo log buffer is written to disk anyway, even if no commit has been sent.

Related Information

[Persistent Data Storage in the SAP HANA Database \[page 462\]](#)

10.2.1.1 Database Restart

An SAP HANA database can be restarted in the same way as a disk-based database, and returned to its most recent consistent state by replaying the redo logs from the log area (not the log backups).

The log entries in the log area only need to be processed from the last savepoint position, rather than from the beginning of the log area. In this way, savepoints help to speed up database restarts.

Backup and Recovery Strategy

While savepoints and logs can protect your data against some failures, these mechanisms offer no protection if the persistent storage itself is damaged or if a logical error occurs. To be able to react appropriately and quickly to a hardware failure, as well as to protect your data against logical errors and the possibility of corruption caused by software changes, it is essential to have a well-planned strategy for backup and recovery.

10.2.2 Points to Note About Backup and Recovery

Before you begin preparing a backup strategy for your SAP HANA installation, you should be aware of specific considerations that apply to backup and recovery.

10.2.2.1 Points to Note: SAP HANA Backups

When you plan your backup strategy, you should be aware of several important points concerning SAP HANA data backups, data snapshots, delta backups, and log backups.

- Backups can only be created when SAP HANA is online.
All the configured SAP HANA services must be running.
While full backups (data backups and data snapshots), delta backups (differential and incremental backups), and log backups are being created, the impact on system performance is negligible, and users can continue to work normally.
For more information, see *SAP HANA Backup Types*.

Caution

Do not create a full backup after a database fault or other failure has occurred.

- With a data backup, only the actual data is backed up; unused space in the database is not backed up. A full data backup includes all the data that is required to recover the database to a consistent state. This includes both business data and administrative data.

Note

A full data backup does not include the log area or customer-specific configuration settings (*.ini files).

- The system database can initiate backups of both the system database itself and of individual tenant databases.

A tenant database can create its own backups without needing to connect through the system database.

- The data backup reflects the consistent database state from the time at which the data backup was started.
Changes made to the database after a data backup was started are not included in the data backup.
If a data backup is recovered without any log backups, open transactions in the data backup are rolled back to the start time of the data backup.
- If a new full backup (data backup or data snapshot) is started before the previous full backup is finished, SAP HANA handles the situation as follows:
 - The first full backup continues normally.
 - The second full backup does not start, and an error message is displayed.
- It is not possible to back up and recover individual database objects.
Backup and recovery always apply to the whole database.
For more information, see *Data Backups*.

SAP HANA Dynamic Tiering and SAP HANA Backup

If you are planning a backup and recovery strategy for a landscape that makes use of SAP HANA dynamic tiering, see SAP Note 2375865 (SAP HANA Dynamic Tiering 2.0: Backup and Restore Functional Restrictions) for information about considerations for dynamic tiering.

i Note

hdbbackupcheck does not support SAP HANA Dynamic Tiering.

Worker Groups

If you have defined worker group sub-roles, information about the worker groups for each volume is stored as part of SAP HANA full backups (complete data backups and data snapshots).

SAP HANA Cockpit and Backup

SAP HANA cockpit 2.0 cannot schedule backups for SAP HANA 1.0 databases.

Related Information

[SAP HANA Backup Types \[page 1246\]](#)

[Data Backups \[page 1246\]](#)

[Tenant Databases \[page 19\]](#)

[Schedule Backups \[page 1325\]](#)

10.2.2.1.1 Points to Note: File-Based Backups

When you plan your backup strategy, you should be aware of several important points concerning how SAP HANA handles backups to the file system.

- The configured destination for data and log backups must be valid throughout the whole system, not only for specific hosts.
Backups of tenant databases are always created in subdirectories of this location.
- To make the backup area available to all the nodes in a database, it is strongly recommended to use shared backup storage.
Shared backup storage allows the system database or the master index server of a tenant database to perform availability checks for file-based backups at the beginning of the recovery.
In addition, shared storage offers support for database copy.

Related Information

[Parameters for Data Backup Settings \[page 1283\]](#)

[Naming Conventions for Backups \[page 1266\]](#)

[Parameters for Backing Up the Backup Catalog \[page 1297\]](#)

10.2.2.1.2 Points to Note: Data Snapshots

Data snapshots offer an additional option to safeguard the SAP HANA data area and to recover an SAP HANA database. If you are planning a backup strategy that makes use of data snapshots, you should be aware of several important points.

- You can create a data snapshot of:
 - An SAP HANA multitenant database container with one tenant database.
Currently, a data snapshot is not supported for an SAP HANA database with more than one tenant database.
To back up SAP HANA systems with more than one tenant database, use data backups.
 - An SAP HANA single-container system
- To **create** a data snapshot, you need to use native SQL.
For more information, see *Create a Data Snapshot (Native SQL)*.
- **Recovery** from a data snapshot is supported by SAP HANA cockpit and SAP HANA studio.

i Note

If you have a backup and recovery strategy that is based on data snapshots, you must ensure that all data snapshots (or at least those you wish to use for a recovery) are replicated outside of the SAP HANA storage system.

Related Information

[Data Snapshots \[page 1249\]](#)

[Create a Data Snapshot \(Native SQL\) \[page 1320\]](#)

[Prerequisites: Recovery From a Data Snapshot \[page 1334\]](#)

[Comparison of Data Backups and Data Snapshots \[page 1395\]](#)

10.2.2.1.3 Points to Note: Log Modes

SAP HANA uses two log modes: `normal` and `overwrite`. By default, SAP HANA runs in log mode `normal`.

After installation, SAP HANA temporarily runs in log mode `overwrite`.

In log mode `overwrite`, no log backups are created.

Log mode `overwrite` ensures that the log area does not grow excessively.

After you create the first full data backup, SAP HANA automatically switches to the default log mode `normal`.

→ Tip

If you change the log mode from `overwrite` – where log backups are not written – to log mode `normal`, **you must create a full data backup** to ensure that log backups are written again, and that the database can be recovered to the most recent point in time.

SAP HANA Dynamic Tiering and Log Mode Overwrite

If you are planning a backup and recovery strategy for a landscape that makes use of SAP HANA dynamic tiering, see SAP Note 2375865 (SAP HANA Dynamic Tiering 2.0: Backup and Restore Functional Restrictions) for information about considerations for dynamic tiering.

Related Information

[Log Modes \[page 1287\]](#)

[Change Log Modes \[page 1290\]](#)

[SAP Note 2375865](#)

10.2.2.1.4 Points to Note: Third-Party Backup Tools

Third-party backup tools can be fully integrated with SAP HANA to enable you to perform backup and recovery operations from SAP HANA studio, SAP HANA cockpit, and using native SQL.

- The implementation of the API of a third-party backup tool that uses the `Backint for SAP HANA` interface must be certified by SAP.
For more information, see *Working with Third-Party Backup Tools*.
- To recover a database, it is possible to use a combination of backups from a third-party backup tool and backups from the file system, provided that the backups originate from the same SAP HANA database. To copy a database, it is not possible to mix backups from the different sources. The backup catalog, the data backups, and the log backups must be from either **only** a third-party backup tool or **only** the file system.
For more information, see *Copying a Database Using Backup and Recovery*.
- SAP HANA supports high isolation scenarios for third-party backup tools.
For more information, see *Isolation Level High for Backups and Third-Party Backup Tools*.

SAP HANA Dynamic Tiering and Third-Party Backup Tools

If you are planning a backup and recovery strategy for a landscape that makes use of SAP HANA dynamic tiering, see SAP Note 2375865 (SAP HANA Dynamic Tiering 2.0: Backup and Restore Functional Restrictions) for information about considerations for dynamic tiering.

Related Information

[Working with Third-Party Backup Tools \[page 1303\]](#)

[Copying a Database Using Backup and Recovery \[page 1374\]](#)

[Isolation Level High for Backups and Third-Party Backup Tools \[page 1308\]](#)

[SAP Note 2375865](#)

10.2.2.1.5 Points to Note: Release Compatibility of SAP HANA Backups

In some situations, backups from earlier SAP HANA releases can be used for a recovery.

- SAP HANA backups created with release 1.0 SPS10 or newer can be used to recover to SAP HANA 2.0. This applies to both SAP HANA single-container systems and tenant databases.
For SAP HANA running on IBM Power systems, different release compatibilities apply.
For more information, see *Points to Note: SAP HANA on IBM Power Systems*.
- A backup of an SAP HANA single-container system can only be recovered to a tenant database.
A backup of an SAP HANA single-container system **cannot be recovered to a system database**.

Related Information

[Points to Note: SAP HANA on IBM Power Systems \[page 1243\]](#)

10.2.2.2 Points to Note: SAP HANA Recovery

Before you plan your backup and recovery strategy, you should be aware of several important points with regard to recovering an SAP HANA database.

- An SAP HANA database cannot be recovered to an SAP HANA database with a lower software version. The software version is the Support Package Stack (SPS) and the database revision. The SAP HANA software version used for the recovery must always be the same or higher than the version of the SAP HANA database used to create the data backup or data snapshot. For more information, see SAP Note 1948334 (*SAP HANA Database Update Paths for Maintenance Revisions*) and SAP Note 2378962 (*SAP HANA 2.0 Revision and Maintenance Strategy*).
- To perform a recovery, an SAP HANA database needs to be shut down. For this reason, during recovery, a database cannot be accessed by end users or applications.

i Note

If you recover SAP HANA from a data snapshot, you must shut down the database **before** you make the data snapshot available in the data area of the storage system.

- To recover a complete SAP HANA system, the system database needs to be recovered first. After the system database has been recovered, each tenant database is recovered individually, and not all together in one single operation. A recovery of a tenant database is always initiated from the system database.
- A system database only needs to be recovered if it is corrupted. If only a tenant database is corrupted, the system database does not need to be recovered.

If the system database is shut down for recovery	All its tenant databases are automatically shut down as well.
---	---

The whole SAP HANA system is not available until the recovery of the system database has been completed.

If a tenant database is shut down for recovery	The system database and any other tenant databases remain online.
---	---

- Using SAP HANA cockpit, only a tenant database can be recovered to a point in time. To recover a system database to a point in time, use SQL. For more information, see *Recovering a Database Using Native SQL*.

SAP HANA Dynamic Tiering

If you are planning a backup and recovery strategy for a landscape that makes use of SAP HANA dynamic tiering, see SAP Note 2375865 (SAP HANA Dynamic Tiering 2.0: Backup and Restore Functional Restrictions) for information about considerations for dynamic tiering.

Worker Groups

If you have defined worker group sub-roles, information about the worker groups for each volume is stored as part of SAP HANA full backups (complete data backups and data snapshots).

Ensure that the target system for a database recovery has the same number of worker groups as the source system.

⚠ Caution

If the target system is not configured correctly, SAP HANA cannot be recovered.

Before you start a recovery, you should normally ensure that the worker groups in the SAP HANA system and the backups have the same names. However, under certain circumstances, it is possible to override this restriction.

For more information, see the `<IGNORE WORKERGROUPS>` option for *RECOVER DATABASE Statement (Backup and Recovery)* in the *SAP HANA SQL and System Views Reference*.

SAP HANA Cockpit and Recovery

SAP HANA cockpit 2.0 SP06 can be used to recover SAP HANA 2.0 and SAP HANA 1.0 databases with Support Package Stack (SPS) 12. For more information about compatible database revisions of SAP HANA 1.0, see 2616241 (*Recovery of SAP HANA 1.0 with SAP HANA Cockpit 2.0*).

SAP HANA cockpit 2.0 SP05 can be used to recover only SAP HANA 2.0 databases.

Related Information

[SAP HANA Recovery \[page 1331\]](#)

[Manually Checking Whether a Recovery is Possible \[page 1335\]](#)

[Points to Note: SAP HANA on IBM Power Systems \[page 1243\]](#)

[Data Temperature: Extension Nodes \[page 581\]](#)

SAP Notes

[SAP Note 1948334](#)

[SAP Note 2378962](#)

[SAP Note 2375865](#)

10.2.2.2.1 Points to Note: Delta Backups and Recovery

SAP HANA supports both differential and incremental backups.

- You can recover an SAP HANA database using a full data backup and a combination of **both** a differential backup and one or more incremental backups.
- By default, when SAP HANA computes a recovery strategy, it gives preference to differential and incremental backups over log backups.
To recover using only a full data backup and log backups, specify the appropriate options in the recovery dialog in SAP HANA cockpit or SAP HANA studio.
- If you recover an SAP HANA database, and do not immediately create a full data backup, the delta backups subsequently created are based on the data backup that was used for the recovery.

i Note

If you wish to recover SAP HANA using differential or incremental backups, you must also use log backups. If log backups are not available, you can only recover using a full data backup.

Related Information

[Delta Backups \[page 1247\]](#)

10.2.2.2.2 Points to Note: License Key and Recovery

When you recover an SAP HANA database, you should be aware of certain license key requirements.

The license key for an SAP HANA database is based on the system ID and the hardware ID. After a recovery, an SAP HANA license key becomes invalid if the SID or hardware ID has changed.

During recovery, a temporary license key is installed automatically if the backup used for recovery had a permanent license, which is still valid. You can work with the automatically installed temporary license for up to 90 days. During this time, you need to apply to SAP to have the license from the source database transferred to a new license key. You then need to install the new license key in the recovered SAP HANA database.

i Note

An SAP HANA license key is installed in the system database. Tenant databases do not need a license key.

⚠ Caution

If the backup that was used for recovery only had a temporary license, the database is in lockdown mode immediately after recovery.

For more information, see *License Keys for SAP HANA Database* in *SAP HANA Administration Guide (Licensing)*.

Related Information

[Recovering an SAP HANA Database \[page 1347\]](#)

[Prerequisites for Copying a Database Using Backup and Recovery \[page 1376\]](#)

[License Keys for the SAP HANA Database \[page 306\]](#)

10.2.2.2.3 Points to Note: Extension Nodes for SAP Business Warehouse

If you are using extension nodes for Business Warehouse, you need to consider some important points with regard to SAP HANA recovery.

Data Temperature

Ensure that warm data is recovered to a service that is configured for warm data.

Likewise, hot data must be recovered to a service that is configured for hot data.

This means that you need to set up the same number and type of services for hot and warm data on the target host that were running on the source host.

For more information, see *Data Temperature: Extension Node for Business Warehouse* and SAP Note 2453736 *How-To: Configuring SAP HANA for SAP BW Extension Node in SAP HANA 2.0*.

Worker Groups

If you have defined worker group sub-roles, information about the worker groups for each volume is stored as part of SAP HANA full backups (complete data backups and data snapshots).

Ensure that the target system for a database recovery has the same number of worker groups as the source system.

Caution

If the target system is not configured correctly, SAP HANA cannot be recovered.

Before you start a recovery, you should normally ensure that the worker groups in the SAP HANA system and the backups have the same names. However, under certain circumstances, it is possible to override this restriction.

For more information, see the `<IGNORE WORKERGROUPS>` option for *RECOVER DATABASE Statement (Backup and Recovery)* in the *SAP HANA SQL and System Views Reference*.

Related Information

[Data Temperature: Extension Nodes \[page 581\]](#)

[SAP Note 2453736](#)

10.2.2.3 Points to Note: Copying a Database Using Backup and Recovery

You can use backup and recovery to copy a tenant database to the same or a different SAP HANA system, or to copy a system database to a different SAP HANA system. When you copy an SAP HANA database, you should be aware of certain important considerations.

The following combinations of source database and target database can be used to create a database copy:

Source Database	Target Database
System database	The system database of a different system
Single-container system	Tenant database
Tenant database	A different tenant database in the same system A tenant database in a different system

i Note

An SAP HANA backup created with SAP HANA 1.0 SPS10 (single-container system) or newer can be used to recover a tenant database.

File System and Third-Party Backup Tools

You can copy an SAP HANA database using file-based backups or backups created using third-party tools.

i Note

To copy a database, it is not possible to mix backups from the different sources. The backup catalog, the data backups, and the log backups must be from either **only** a third-party backup tool or **only** the file system.

(To recover a database, it is possible to use a combination of backups from a third-party backup tool and backups from the file system, provided that the backups originate from the same SAP HANA database.)

Database Copy and Data Snapshots

Currently, it is only possible to use a data snapshot to back up and recover an SAP HANA single-tenant system.

To recover SAP HANA from a data snapshot, you first need to recover the system database, then the tenant database.

i Note

If you recover SAP HANA from a data snapshot, you must shut down the database **before** you make the data snapshot available in the data area of the storage system.

⚠ Caution

It is **not possible** to use a data snapshot of an SAP HANA single-container system to recover an SAP HANA multitenant database container.

Database Copy and System Replication

If you have system replication configured, and require near-zero downtime, consider using system replication to copy a tenant database.

For more information, see *Copying and Moving Tenant Databases Between Systems* in the *SAP HANA Administration Guide*.

Database Copy and Delta Backups

If you wish to create a database copy using differential or incremental backups, you must also use log backups. If log backups are not available, you can only create a database copy using a full data backup.

Related Information

[Copying a Database Using Backup and Recovery \[page 1374\]](#)

[Copying and Moving Tenant Databases Between Systems \[page 1004\]](#)

10.2.2.4 Points to Note: System Replication

Data backups and log backups can only be written on the primary system.

After a Takeover

i Note

After a takeover, it is not necessary to create a new full data backup (data backup or data snapshot) of the now active system. Backups of the former primary system can be used to recover the database.

However, no delta backups can be created in the now active system until a full data backup has been created.

If you wish to recover SAP HANA using differential or incremental backups, you must also use log backups. If log backups are not available, you can only recover using a full data backup.

After a takeover, ensure the following:

- Backups from the former primary system are not being written to the same location as backups from the now active system.

⚠ Caution

If backups from different systems are mixed up, it will not be possible to recover the database.

This can be achieved in either of the following ways:

- Deactivate automatic log backups and any scheduled data backups in the former primary system.
- Shut down the former primary system to ensure that it creates no new data backups and no new log backups.
- Any backups scheduled in the now active system are configured in accordance with your requirements. For more information, see *Schedule Backups*.

Before a Recovery

Before a recovery, disable the FULL SYNC option.

If you are running system replication with replication mode **SYNC** and the **FULL SYNC** option enabled, the system will not start after a recovery, because no write operations are possible.

To prevent this from happening, before you perform a recovery, manually disable the **FULL SYNC** option in `global.ini`.

You can use the following command as `<sid>adm`:

```
hdbnsutil -sr_fullsync --disable
```

For more information, see SAP Note 2165547 (FAQ: SAP HANA Database Backup & Recovery in an SAP HANA System Replication Landscape) and *Recovery with System Replication*.

System Replication and Third-Party Backup Tools

- If backups are managed using a third-party tool, the `Backint for SAP HANA` API must be accessed by both the active system and the original primary system.
- If SAP HANA is recovered from backups that were created with different UIDs, some third-party backup tools may prevent the recovery from being started.
For more information, contact your tool vendor or ensure that the same UID is used for all the backups used for a recovery.

Related Information

[Schedule Backups \[page 1325\]](#)

[Recovery with System Replication \[page 1371\]](#)

[SAP Note 2165547](#)

10.2.2.5 Points to Note: SAP HANA on IBM Power Systems

If you are working with IBM Power systems, you should be aware of certain important points concerning SAP HANA.

- Backups created with SAP HANA 2.0 are compatible with the supported hardware platforms Intel and IBM Power. You can recover SAP HANA 2.0 using data backups and log backups created with SAP HANA 2.0 on either an Intel-based system or an IBM Power-based system.
Data backups and log backups created with SAP HANA 1.0 SPS10 or newer running on an Intel-based system can be used to recover SAP HANA 2.0 to both Intel-based and IBM Power-based systems. Data backups and log backups created with SAP HANA 1.0 on an IBM Power-based system **cannot** be used to recover SAP HANA 2.0.

Compatibility of Backups of SAP HANA 1.0 for Recovery to SAP HANA 2.0 (IBM Power and Intel)

Backup Source (Data Backups and Log Backups)	Recovery to SAP HANA 2.0 (IBM Power)	Recovery to SAP HANA 2.0 (Intel)
	SAP HANA 1.0 SPS9 and earlier	NO
SAP HANA 1.0 (SPS10 and later) (IBM Power)	NO	NO
SAP HANA 1.0 (SPS10 and later) (Intel)	YES	YES
SAP HANA 2.0 (IBM Power and Intel)	YES	YES

- For third-party backup tools, separate certification processes are required for each platform and tool version.

If a third-party backup tool is certified for Intel platforms, that tool is **not** automatically also certified for IBM Power Systems (and vice versa).

Separate tool certification is required for IBM Power LE and IBM Power BE systems.

10.2.3 Authorizations for Backup and Recovery

Backup and recovery operations can only be performed by users that have the appropriate authorizations. The authorization required depends on whether administrative tasks are performed at system level or at database level.

The following authorizations are required to administer SAP HANA:

Authorizations for Backup and Recovery

Task / SAP Tool	System Database	Tenant Database	Tenant Database
			(Through the System Database)
Backup (SAP HANA Cockpit)	BACKUP ADMIN or BACKUP OPERATOR (recommended for batch users only)	BACKUP ADMIN or BACKUP OPERATOR (recommended for batch users only)	DATABASE ADMIN
Backup (SAP HANA Studio)	BACKUP ADMIN CATALOG READ	BACKUP ADMIN CATALOG READ	DATABASE ADMIN
Backup (Native SQL)	BACKUP ADMIN or BACKUP OPERATOR (recommended for batch users only)	BACKUP ADMIN or BACKUP OPERATOR (recommended for batch users only)	DATABASE ADMIN
Recovery (SAP HANA Cockpit)	Operating system user <sid>adm	(not possible)	DATABASE ADMIN
Recovery (SAP HANA Studio)	Operating system user <sid>adm	(not possible)	DATABASE ADMIN
Recovery (Native SQL)	Operating system user <sid>adm	(not possible)	DATABASE ADMIN
Schedule Backup (SAP HANA Cockpit)	BACKUP ADMIN	BACKUP ADMIN	(Currently not supported)
Administration Tasks	BACKUP ADMIN	BACKUP ADMIN	DATABASE ADMIN

For example, physically delete data backups, log backups, and obsolete versions of the backup catalog.

→ Tip

We recommend that you create your own dedicated database users with only the specific authorizations required for backup and recovery.

BACKUP ADMIN Versus BACKUP OPERATOR

The system privileges BACKUP ADMIN and BACKUP OPERATOR allow you to implement a more specific separation of user roles.

With BACKUP ADMIN, a user can perform **all** backup-related operations, including deleting backups and backup configurations. With BACKUP OPERATOR, a user can only perform backups.

For example, if you have automated the regular execution of backups using cron, it is more secure to use a user with the authorization BACKUP OPERATOR, as this prevents backups from being deleted inadvertently.

For more information, see *SAP HANA Authorization* in the *SAP HANA Security Guide*.

Related Information

[Operating System User sidadm \[page 714\]](#)

10.2.4 SAP HANA Backup

There are different methods and tools to back up an SAP HANA database.

The following sections describe:

- The backup types supported by SAP HANA
- Redo log backups
- Naming conventions
- Backing up customer-specific configuration settings
- The backup catalog
- Backup storage in the file system and using third-party backup tools
- Creating an SAP HANA backup
- Backup audit actions for security
- Checking the integrity of backups

Related Information

[SAP HANA Backup Types \[page 1246\]](#)

[Log Backups \[page 1252\]](#)

[Naming Conventions for Backups \[page 1266\]](#)

[Backing Up Customer-Specific Configuration Settings \[page 1264\]](#)

[Backup Catalog \[page 1255\]](#)
[Persistent Data Storage in the SAP HANA Database \[page 462\]](#)
[Working with Third-Party Backup Tools \[page 1303\]](#)
[Creating Backups \[page 1313\]](#)
[Backup Audit Actions for Security \[page 1331\]](#)
[Checking Individual Backups \[page 1337\]](#)

10.2.4.1 SAP HANA Backup Types

SAP HANA supports different backup types.

Related Information

[Data Backups \[page 1246\]](#)
[Delta Backups \[page 1247\]](#)
[Data Snapshots \[page 1249\]](#)
[Log Backups \[page 1252\]](#)
[SAP HANA Backup Encryption \[page 1252\]](#)
[Backup Catalog \[page 1255\]](#)

10.2.4.1.1 Data Backups

A data backup includes all the data that is required to recover the database to a consistent state.

With a data backup, only the actual data is backed up; unused space in the database is not backed up.

i Note

A data backup does not include the log area or customer-specific configuration settings (*.ini files).

The data area is backed up in parallel for each of the SAP HANA services. If SAP HANA is running on multiple hosts, a data backup includes all the service-specific backup parts for all the hosts.

While a data backup is running, some data integrity checks are performed. These integrity checks are performed on block level (page level on disk) only, and do not analyze the content of each data block.

If these checks are successful, the data is written to the backup destination.

i Note

To ensure the safety of your data, data backups should be stored on multiple different backup destinations outside the SAP HANA database.

Related Information

[Creating Backups \[page 1313\]](#)

[Backing Up Customer-Specific Configuration Settings \[page 1264\]](#)

10.2.4.1.2 Delta Backups

Delta backups back up only data that has been changed since the last full data backup (complete data backup or data snapshot) or the last delta backup.

i Note

Delta backups can only be created after a data backup has been created.

With a delta backup, changed data means changes to the physical representation of the data in the SAP HANA persistent storage. This is not always the data that was actually changed by an application. An internal reorganization can change the physical representation **without changing the actual data**.

🔗 Example

In a delta merge of a column store partition, only a small amount of the data may have been changed. Nevertheless, all the data is restructured and rewritten. This means that a delta merge can be as large as a full data backup. If a delta backup is created in this situation, the whole partition is backed up in the delta data backup, even if only a small amount of the actual data was changed.

Benefits of Delta Backups

Delta backups allow you to reduce the amount of data that is backed up, compared to full data backups.

In turn, this means that delta backups are normally faster to create than full data backups.

In addition, a database recovery using delta backups is normally faster than a recovery using log backups. With delta backups, only the changed data is recovered, whereas with log backups, each log entry needs to be processed separately before it is recovered. Recovering many log backups is normally more CPU-intensive than recovering a small number of delta backups.

→ Tip

To keep data that is frequently changed separate from data that is not frequently changed, consider partitioning column store data. Partitioning column store data can reduce the size of delta backups.

Related Information

10.2.4.1.2.1 Delta Backup Types

SAP HANA supports both differential backups and incremental backups.

i Note

Delta backups can only be created after a data backup has been created.

Comparison of Delta Backup Types

	Differential Backup	Incremental Backup
What Data is Backed Up?	The data changed since the last full data backup.	The data changed since the last full data backup or the last delta backup (incremental or differential).
Backup Size	The amount of data to be saved with each differential backup increases .	If data remains unchanged, it is not saved to more than one backup. For this reason, incremental backups are the smallest of the backup types.
Backup and Recovery Strategy	If your backup strategy is based on only full data backups and differential backups, only two backups are needed for a recovery: one full data backup and one differential backup.	<p>If your backup strategy is based on only full data backups and incremental backups, to recover the database, SAP HANA needs the following backups:</p> <ul style="list-style-type: none"> • The full data backup on which the incremental backups are based • Each incremental backup made since the full data backup <p>In some situations, many incremental backups may be needed for a recovery.</p>

Recovery Using Delta Backups

i Note

A recovery can use multiple incremental backups, but only one differential backup.

If you wish to recover SAP HANA using differential or incremental backups, you must also use log backups. If log backups are not available, you can only recover using a full data backup.

To recover SAP HANA, you can combine a differential backup with one or more incremental backups.

❖ Example

You could recover SAP HANA to a specific point in time using the following sequence of backups:

1. Full data backup
2. Differential backup
3. Incremental backup 1
4. Incremental backup 2
5. Log backups

10.2.4.1.3 Data Snapshots

A data snapshot captures the data persisted in the data area at a particular point in time. A data snapshot includes all the data that is required to recover SAP HANA to a consistent state.

i Note

To make use of storage snapshot-based SAP HANA backups, first create an SAP HANA data snapshot.

- You can create a data snapshot of:
 - An SAP HANA multitenant database container with one tenant database.
Currently, a data snapshot is not supported for an SAP HANA database with more than one tenant database.
To back up SAP HANA systems with more than one tenant database, use data backups.
 - An SAP HANA single-container system
- To **create** a data snapshot, you need to use native SQL.
For more information, see *Create a Data Snapshot (Native SQL)*.
- **Recovery** from a data snapshot is supported by SAP HANA cockpit and SAP HANA studio.

Benefits of Data Snapshots

Data snapshots offer an additional option to safeguard the SAP HANA data area and to recover an SAP HANA database.

Data snapshots have the following benefits:

- Data snapshots can be created with minimal impact on the system.
This is because data snapshots are created in the storage system and do not consume database resources.
- Recovery from a data snapshot is faster than a recovery from a data backup.
The data snapshot only needs to be made available in the data area of the storage system.

For more information about the relative benefits of data snapshots, see *Comparison of Data Backups and Data Snapshots*.

Data Snapshots and Database Copy

For a database copy using data snapshots, the number of hosts and the number and type of services assigned to each host must be the same for the source database and the target database, and the mountpoint IDs must be identical.

For more information, see *Prerequisites for Copying a Database Using Backup and Recovery*.

Data Snapshots and External Storage Systems

The external storage system must copy each data volume in an atomic operation in order to ensure the I/O consistency of the data snapshot. Multiple data volumes do not need to be copied in parallel; data volumes can be copied one at a time.

Data Snapshots and SAP HANA Dynamic Tiering

If you are planning a backup and recovery strategy for a landscape that makes use of SAP HANA dynamic tiering, see SAP Note 2375865 (SAP HANA Dynamic Tiering 2.0: Backup and Restore Functional Restrictions) for information about considerations for dynamic tiering.

Related Information

[Comparison of Data Backups and Data Snapshots \[page 1395\]](#)

[Create a Data Snapshot \(Native SQL\) \[page 1320\]](#)

[Encryption of Data Snapshots \[page 1251\]](#)

[SAP Note 2375865](#)

[Prerequisites: Recovery From a Data Snapshot \[page 1334\]](#)

[Prerequisites for Copying a Database Using Backup and Recovery \[page 1376\]](#)

10.2.4.1.3.1 Data Snapshots and Database Snapshots

A **data snapshot** is created by first creating an **internal database snapshot**. This internal database snapshot provides a view of the database at the point in time that it was started.

The internal database snapshot is used to ensure the consistent state of the data snapshot.

i Note

The internal database snapshot reflects a consistent state. While a **data snapshot** is being created (based on the **internal database snapshot**), no further data integrity checks are performed.

(With data backups, the block-level integrity of the data to be backed up is checked automatically while the backups are being created.)

Internal Database Snapshot and System Replication

An internal database snapshot used to create a data snapshot does not conflict with an internal database snapshot used for system replication. There is no relation between these two types of internal database snapshots.

Related Information

[Create a Data Snapshot \(Native SQL\) \[page 1320\]](#)

10.2.4.1.3.2 Encryption of Data Snapshots

Encryption of Data Snapshots

A data snapshot can be encrypted if the following conditions are met:

- Backup encryption is enabled.
- Data volume encryption is enabled.
For more information, see the *SAP HANA Administration Guide (Encryption)* and the *SAP HANA Security Guide*.
- The data conversion status is **not active**.
The SAP HANA database must not be in the process of encrypting or decrypting data.
To check the data conversion status, execute the statement:

```
SELECT DATA_CONVERSION_ACTIVE FROM M_PERSISTENCE_ENCRYPTION_STATUS
```


To encrypt a data snapshot, `data_conversion_active` must be **false**.
For more information, see *M_ENCRYPTION_OVERVIEW System View* in the *SAP HANA SQL and System Views Reference*.

Related Information

[SAP HANA Backup Encryption \[page 1252\]](#)

10.2.4.1.4 Log Backups

By default, SAP HANA log segments in the log area are backed up automatically.

Log segments are backed up for each service that has persistence. During a log backup, only the actual data (the "payload") of the log segments is written from the log area to service-specific log backups in the file system or to a third-party backup tool.

⚠ Caution

If an SAP HANA service stops, log backups for that service also stop. The stopped service must be immediately restarted.

If the stopped service is not restarted, a database recovery will only be possible to a point in time before this service stopped. That is, only to a point in time for which log backups for **all services** exist.

If log backups for any service are missing, it will not be possible to recover the database to its most recent state.

i Note

If the log backup area becomes temporarily unavailable, once it is available again, SAP HANA automatically continues creating log backups for all the log segments that have not so far been backed up.

Removing a Service

If you need to remove a service, use the procedure described in the section *Steps After Copying a Database*, as this ensures that the log area is backed up and can be used to recover the database.

Related Information

[Log Modes \[page 1287\]](#)

[Naming Conventions for Log Backups \[page 1270\]](#)

[Backup Configuration Settings \[page 1275\]](#)

[Steps After Copying a Database \[page 1389\]](#)

10.2.4.1.5 SAP HANA Backup Encryption

SAP HANA supports native encryption of backups.

Backup encryption safeguards the privacy of the SAP HANA business data by preventing unauthorized parties from reading the content of backups.

i Note

As an alternative to native SAP HANA encryption, many third-party backup tools and storage tools offer support for backup encryption. If you are using a third-party backup tool, consult your tool vendor for more information.

Which Backup Types Can Be Encrypted?

Backup encryption can be enabled for all backup types.

For more information, see *Enable and Disable Encryption of Data and Log Backups* in *SAP HANA Administration Guide (Encryption)*.

i Note

To encrypt data snapshots, additional steps are necessary.

For more information, see *Encryption of Data Snapshots* in *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

Considerations for Backup Encryption

- It takes longer to create encrypted backups than unencrypted backups.
- It takes longer to recover a database using encrypted backups than from unencrypted backups.
- If backup encryption is enabled, both data backups and log backups are encrypted.

i Note

If you enable encryption (either for backup, log, or data volume) in the system database immediately after installation of SAP HANA, encryption is automatically enabled for any subsequently created tenant databases. If a particular tenant database does not require encryption, the administrator for that tenant database can disable encryption for it.

For more information, see *Enable Data and Log Volume Encryption in a New SAP HANA Database* in *SAP HANA Administration Guide (Encryption)*.

- The same backup encryption root key is used for both data backups and log backups.

i Note

Data snapshots are not encrypted using the backup encryption root key.

For more information, see *Encryption of Data Snapshots* in *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

- It is currently not possible to enable encryption for an individual data backup.
- The size of encrypted SAP HANA backups is the same as unencrypted backups (except for the checksum).
- The backup catalog is not encrypted.
The backup catalog shows which backup encryption root keys were used to encrypt the backups.

Backup Encryption Root Keys

- If you enable encryption (either for backup, log, or data volume) in the SAP HANA backups are encrypted and decrypted using backup encryption root keys.

The backup encryption root keys are encrypted and stored in the secure store in the file system (instance SSFS) together with other encryption root keys. For example, application encryption root keys.

- A new backup encryption root key is generated for every tenant database when the tenant database is created.
- If backup encryption is enabled, a database administrator must ensure that the backup encryption root keys are backed up.

⚠ Caution

Whenever the backup encryption root keys are changed, you must back them up. Without a current backup of the backup encryption root keys, some data will be lost after a recovery.

For more information, see *Root Key Backup* in the *SAP HANA Security Guide*.

i Note

The block-level integrity of encrypted backups can still be checked without access to the backup encryption root keys.

For more information, see *Manually Checking Whether a Recovery is Possible*.

For more information about working with encryption root keys, see *Encryption Key Management* in the *SAP HANA Security Guide* and *Changing Encryption Root Keys* in the *SAP HANA Administration Guide (Encryption)*.

SAP HANA Dynamic Tiering and Backup Encryption

If you are planning a backup and recovery strategy for a landscape that makes use of SAP HANA dynamic tiering, see SAP Note 2375865 (SAP HANA Dynamic Tiering 2.0: Backup and Restore Functional Restrictions) for information about considerations for dynamic tiering.

Related Information

SAP HANA Administration - Database Backup and Recovery

[Encryption of Data Snapshots \[page 1251\]](#)

[Manually Checking Whether a Recovery is Possible \[page 1335\]](#)

SAP HANA Security Guide

SAP HANA Administration - Security Administration

[Import Backed-up Root Keys \[page 876\]](#)

[Changing Encryption Root Keys \[page 858\]](#)

[Server-Side Data Encryption Services \[page 848\]](#)

[Enabling Encryption of Data and Log Volumes \[page 864\]](#)

[Enable Data and Log Volume Encryption in a New SAP HANA Database \[page 865\]](#)

[Enable Encryption of Data and Log Backups \[page 872\]](#)

10.2.4.1.6 Backup Catalog

The backup catalog contains information about the backup history.

The backup catalog enables SAP HANA to determine the following:

- Whether a recovery is possible
- Which backups to use to recover a database
- Which backups are no longer needed for a recovery

The system database and each tenant database have their own backup catalog.

i Note

Each time a backup of any type is created, the backup catalog is backed up and versioned. In this way, the latest version of the backup catalog always contains the entire backup history.

Recovery Without the Backup Catalog

It is possible to recover SAP HANA without using a backup catalog, or using data backups that are not recorded in the backup catalog.

Without a backup catalog, a point-in-time recovery cannot be done. It is only possible to recover SAP HANA to a specific data backup.

If you recover SAP HANA using a backup that is not recorded in the backup catalog, you need to manually specify the backup type (File or Backint), the location of the backup, and its prefix.

Related Information

Working with the Backup Catalog

[What Information is in the Backup Catalog? \[page 1256\]](#)

[Monitoring Views for the Backup Catalog \[page 1258\]](#)

[Housekeeping for Backup Catalog and Backup Storage \[page 1259\]](#)

[Rebuilding the Backup Catalog \[page 1263\]](#)

Backup Catalog Configuration

[Backup Configuration Settings \[page 1275\]](#)

[Naming Conventions for the Backup Catalog \[page 1271\]](#)

[Configure Backups \[page 1274\]](#)

10.2.4.1.6.1 What Information is in the Backup Catalog?

The backup catalog contains information about the backup history.

In the Backup Catalog	Description
Backups created for a database	<p>This includes data backups, data snapshots, delta backups (differential and incremental backups), and log backups.</p> <p>Each recovery is recorded in the backup catalog, but not displayed in the monitoring views.</p> <p>For more information, see <i>Monitoring Views for the Backup Catalog</i>.</p>
Start and completion times of the backups	<p>The backup catalog records local server times and UTC times.</p> <p>The start and end times of a recovery are recorded in the backup catalog as local server time. Points in time for a recovery in the SQL statements are specified as UTC.</p> <p>The start time of a backup does not reflect the exact state of the database when the backup was created. The savepoint that determines the database state is always taken after the start time of a data backup.</p> <p>Changes made to the database after this savepoint are not included in a data backup.</p>
Whether a backup is still running	<p>The backup catalog does not show the progress of a backup. The progress of a backup is recorded in the <code>backup.log</code>, and can be seen using the view <code>M_BACKUP_PROGRESS</code>.</p> <p>For more information, see <i>backup.log</i> and <i>Monitoring Views for the Backup Catalog</i>.</p>
Status of a backup	<p>Records whether a backup was completed successfully or not.</p> <p>The status can be:</p> <ul style="list-style-type: none">• successful• failed• running• cancel pending• canceled
Volumes that were backed up	<p>Backups are created in separate backup destinations for each volume.</p>

In the Backup Catalog	Description
Log backups and the log positions they contain	<p>For data backups, one exact redo log position is recorded. This redo log position corresponds to its savepoint. During a recovery, log replay starts at the redo position for the data backup.</p> <p>For log backups, a range of redo log positions is recorded, from the oldest log entry to the newest log entry.</p>
Backup destinations and their sizes	<p>Data backups can be written to different destinations. However, all the parts of the same data backup are written to the same destination.</p> <p>For file-based data backups, you can change the default destination.</p> <p>For backups created using third-party backup tools, a named pipe is created in the file system for the system database or tenant database. The default backup destination cannot be changed.</p>
Destination type	The destination type can be: file, Backint, or data snapshot
Backup ID	<p>ID of a data backup or a log backup.</p> <p>All backup files of a single data backup share the same BACKUP_ID.</p> <p>If you are working with a third-party backup tool, an external backup ID (EBID) is also included.</p>
Encryption Root Key	If encryption is enabled, the key used for encrypting a backup.
Comment	<p>Additional user-supplied information.</p> <p>A comment can help to identify a particular backup in the backup catalog.</p>

Related Information

[Monitoring Views for the Backup Catalog \[page 1258\]](#)
[backup.log \[page 1272\]](#)

10.2.4.1.6.2 Monitoring Views for the Backup Catalog

You can use monitoring views to display information from the backup catalog. Monitoring views are stored in the SYS schema.

The monitoring views `M_BACKUP_CATALOG`, `M_BACKUP_CATALOG_FILES`, and `M_BACKUP_PROGRESS` provide different overviews of information from the backup catalog.

Monitoring View	Description
<code>M_BACKUP_CATALOG</code>	<p>Provides an overview of information about backup activities.</p> <p>Each row in the view provides information about a separate catalog entry identified by a backup ID. This information includes the type (for example, data backup), and start and completion times.</p> <div style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>The backup ID is used to reference the parts of a backup in the <code>M_BACKUP_CATALOG_FILES</code> monitoring view.</p> </div>
<code>M_BACKUP_CATALOG_FILES</code>	<p>Provides information about the backups created, and the backup destinations for data and log backups.</p> <p>Each row in the view has a corresponding entry in the <code>M_BACKUP_CATALOG</code> monitoring view. Each row is identified by a backup ID.</p> <p>The <code>M_BACKUP_CATALOG_FILES</code> monitoring view provides additional information about each database service that was involved in a backup. For example, with a data backup, each database service is listed with its specific backup information such as destination path and redo log position.</p>
<code>M_BACKUP_PROGRESS</code>	<p>Provides detailed information about the most recent data backup.</p> <p>Each row contains information about one service that is part of the data backup, identified by host name and port number.</p>

Comparison of the Monitoring Views for the Backup Catalog

<code>M_BACKUP_CATALOG</code>	<code>M_BACKUP_CATALOG_FILES</code>	<code>M_BACKUP_PROGRESS</code>
All types of backups (data backup, log backup, delta backups, and data snapshots, if available)		Only for data backups, differential backups, and incremental backups
All completed and currently running backups since the database was created		Currently running and last finished backups only
Persistent		Cleared at database restart
Total amount of data for completed backups only		Total and already transferred amount of data for all backups

For more information, see the *System Views Reference*.

❖ Example

Search for a Data Snapshot Using M_BACKUP_CATALOG

To search for a data snapshot in the backup catalog, you can use either the backup ID or the comment.

To search for a backup ID, use the following command:

```
SELECT * FROM "SYS"."M_BACKUP_CATALOG" WHERE BACKUP_ID = backup_id
```

10.2.4.1.6.3 Housekeeping for Backup Catalog and Backup Storage

It is recommended to regularly delete backups that are no longer needed.

You can physically delete data backups and log backups, and delete their associated records in the backup catalog. For example, you can delete individual data backups if they can no longer be used for a recovery, or to keep your backup storage space at an optimum level.

You can also delete the records of individual data backups from the backup catalog, but retain the physical backups. For example, to comply with legal requirements for data retention, you may wish to retain specific historical data backups, but without the intention of using them for a database recovery.

→ Tip

It is important to regularly truncate the backup catalog because, as it increases in size, it can consume a lot of storage space and also take longer to write each new backup.

To delete data backups, it is recommended to use the function *Delete Older Backups...* in the Backup Console in SAP HANA studio, as this also deletes log backups and delta backups that are no longer needed for a recovery, while ensuring that SAP HANA can still be recovered to a consistent state.

i Note

If a data backup is physically available, but not recorded in the backup catalog, that data backup can still be used to recover the database.

Third-party data backups that are not recorded in the backup catalog must have a unique prefix to be able to be used for a recovery.

Related Information

[Backup Catalog \[page 1255\]](#)

10.2.4.1.6.3.1 Prerequisites for Deleting Old Backups

Before you delete old data backups, delta backups, and log backups, some prerequisites must be met.

- You have the appropriate system privilege:

BACKUP ADMIN

Connect directly to a system database or a tenant database and delete backups of that database.

To connect directly to a tenant database to delete old backups, use native SQL.

For more information, see *BACKUP CATALOG DELETE Statement (Backup and Recovery)*.

DATABASE ADMIN

Connect to a system database and delete backups of a tenant database.

- To delete backups of a tenant database, the tenant database must be online.
- To delete multiple backups, you need to decide from which time onwards you want to retain data backups.
- Before you physically delete backups, ensure that the versions retained are accessible and consistent.
- At least one data backup must remain in the backup catalog.

Deleting Data Snapshots

If data snapshots are deleted from the backup catalog, they are **not deleted physically**.

Data snapshots that are no longer needed must be deleted manually.

Deleting File System Backups

SAP HANA searches for a backup only in the physical location recorded in the backup catalog. If a backup has been moved from the location recorded in the backup catalog, SAP HANA cannot delete it.

Archiving Backups

Backups that need to be retained for extended periods can be removed from the backup catalog and archived in a secure location. Ensure that these backups cannot be accessed directly by SAP HANA and cannot be deleted.

An archived **data backup** can still be used to recover SAP HANA, even if it is not recorded in the backup catalog. If you need to ensure that you can recover SAP HANA from older log backups and delta backups, you need to retain backups of the backup catalog.

Related Information

[Parameters for Backing Up the Backup Catalog \[page 1297\]](#)

10.2.4.1.6.3.2 Delete Old Backups

You can delete records of backups from the backup catalog. Alternatively, you can delete records of backups from the backup catalog and also delete the associated physical backups.

Context

⚠ Caution

Deleting backups is a critical operation. To keep a record of all backup deletions, use the auditing feature in the SAP HANA database. The action to be audited is BACKUP CATALOG DELETE.

For more information about auditing, see *Auditing Activity in the SAP HANA Database* and *Create an Audit Policy* in the *SAP HANA Administration Guide*.

Procedure

i Note

The steps described here ensure that no backups are deleted that would prevent SAP HANA from being recovered. For this reason, it is not possible to delete delta backups and log backups individually.

1. In SAP HANA studio, choose **► Backup and Recovery ► Open Backup Console**.
2. In the Backup Console for the system database, go to the **Backup Catalog** tab.
3. Select the database.
4. Select a data backup and right-click to open the context menu.

Option	Description
Delete Data Backup...	Delete the selected data backup.
Delete Older Backups...	Retain the selected data backup and delete all older data backups, delta backups, and log backups.

A dialog box appears.

5. Specify what you want to delete:

Option	Description
<i>Catalog</i>	Delete the records of the selected data backup(s) from the backup catalog only.
<i>Catalog and Backup Location</i>	<p>Delete the selected data backup(s) from the backup catalog and also delete the physical backups.</p> <p>Specify the location of the physical data backup(s) to be deleted.</p> <p>If you have physical backups in both the file system and a third-party backup tool, you can choose to delete data backups in only one location.</p>

i Note

SAP HANA can only physically delete backups that are in the location recorded in the backup catalog. SAP HANA cannot physically delete backups that have been moved to a different location.

6. Choose *Next* and review the information about the data backup(s) that will be deleted.
Optionally, you can download a list of the deleted backups to a plain text file.
7. If you are sure that you want to delete the data backup(s), choose *Finish*.

Results

The system deletes the selected data backup(s).

If you chose to delete from the backup catalog only, the system deletes the backup records, and refreshes the backup catalog in the Backup Console.

i Note

If you chose to delete records from the backup catalog **and also the associated physical backups**, the truncated backup catalog is displayed immediately, even though it may take some more time for the physical backups to be deleted.

Before the backup(s) are physically deleted, the following plausibility checks are performed:

- **For file-based backups:**
The system checks that the backup ID matches the landscape ID of the current database.
- **For third-party backup tools:**
The system checks that the path to the backup is identical to the backup destination of the current database.

If the plausibility check is successful, the system starts deleting the physical backup(s).

Depending on the size and number of the backups, this can take some time.

You can monitor the progress of the deletion operation in the `backup.log` file.

i Note

The delete operation continues until all the selected backups have been deleted.

If the system or a service is stopped and restarted, the delete operation is automatically resumed.

If the oldest data backup is deleted, the delta backups and the log backups for the period up to the next oldest data backup are **not** automatically deleted.

Related Information

[Auditing Activity in the SAP HANA Database \[page 826\]](#)

[Create an Audit Policy \[page 839\]](#)

[Manually Checking Whether a Recovery is Possible \[page 1335\]](#)

[backup.log \[page 1272\]](#)

10.2.4.1.6.4 Rebuilding the Backup Catalog

In exceptional situations that are outside of the control of SAP HANA, the backup catalog may not be available at the time of a recovery. If the backup catalog is not available, it can be largely rebuilt using the existing data and log backups from the file system.

To rebuild the backup catalog, use the `hdbbackupdiag` tool.

For more information, see SAP Note 1812057 (*Reconstruction of the backup catalog using hdbbackupdiag*).

⚠ Caution

If the backup catalog is rebuilt:

- The backup catalog only contains information about the data backups and log backups that you provide through the `hdbbackupdiag` tool.
Any parts of the database backup history that you do not specify is lost.
- The backup catalog no longer contains backups made using third-party tools or data snapshots.
As a consequence, those backups cannot be used for a recovery.
- Data and log backup directories must contain **only SAP HANA data**.
- A log is generated and written to the SAP HANA working directories.
If you rebuild the backup catalog **more than once**, this log will also be read, and an error will be caused.
For this reason, if you need to rebuild the backup catalog a second time, remove the log from the first rebuild.
- Only the backup of the backup catalog is affected; the persistent SAP HANA backup catalog is not changed.
The backup of the backup catalog is used to recover SAP HANA.

Related Information

[SAP Note 1812057](#) 

[Manually Checking Whether a Recovery is Possible \[page 1335\]](#)

10.2.4.1.7 Backing Up Customer-Specific Configuration Settings

Customer-specific configuration settings (changes to the *.ini files) are not backed up automatically as part of a full backup. The configuration settings are not essential to perform a database recovery. If you want to back up configuration files that contain customer-specific changes, you can do so manually.

In a recovery situation, a backup of the configuration settings can be helpful to more easily identify and restore customer-specific changes to the default settings. If you want to use customer-specific configuration settings after a recovery, you need to reconfigure the recovered system.

You can display and change configuration parameters using SAP HANA cockpit and SAP HANA studio.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

Locations of the SAP HANA Configuration Files

By default, configuration files for SAP HANA are written to specific directories.

Locations of the SAP HANA Configuration Files

Configuration Settings	Location
Global configuration settings	<code>/usr/sap/\$SID/global/hdb/custom/config/</code>
Configuration settings for a tenant database	<code>/usr/sap/\$SID/global/hdb/custom/config/ DB_<database_name></code> For more information, see <i>Database-Specific Configuration Parameters</i> in the <i>SAP HANA Administration Guide</i> .
Host-specific configuration settings	<code>/usr/sap/<SID>/HDB<instance no.>/<host>/</code>

i Note

Configuration files (.ini files) are only created if customer-specific changes are made to them after installation. If no customer-specific changes have been made, the directories may be empty.

Default SAP HANA Configuration Files

During installation of SAP HANA database, the following configuration files are created:

Content of the Main SAP HANA Configuration Files

Configuration File	Description
<code>sapprofile.ini</code>	<p>Contains system identification information, such as the system name (SID) or the instance number.</p> <p>After installation, <code>sapprofile.ini</code> is not changed again.</p> <div data-bbox="821 683 1396 913"><p>⚠ Caution</p><p><code>sapprofile.ini</code> contains information that is specific to each host. For this reason, in a recovery situation, the <code>sapprofile.ini</code> file must not be copied manually to a different host, as it will not be compatible with a new landscape.</p></div> <p>The <code>sapprofile.ini</code> file can be found in the following directory:</p> <pre>/usr/sap/<SID>/HDB<instance no.>/<host>/</pre> <div data-bbox="821 1070 1396 1272"><p>i Note</p><p><code>sapprofile.ini</code> is not displayed in the <i>Configuration of System Properties</i> overview in SAP HANA cockpit or in the <i>Configuration</i> tab in SAP HANA studio.</p></div>
<code>daemon.ini</code>	<p>Contains information about which database services to start.</p> <p>The <code>daemon.ini</code> file can be found in the following directory:</p> <pre>/usr/sap/<SID>/HDB<instance no.>/<host>/</pre>
<code>nameserver.ini</code>	<p>The <code>nameserver.ini</code> file contains global configuration settings for each SAP HANA installation.</p> <p>The landscape section contains the system-specific landscape ID and assignments of hosts to roles MASTER, WORKER, and STANDBY. It also contains configuration settings for system replication and for the SAP HANA Storage Connector API.</p> <p>If the system landscape is changed, for example, hosts are added or removed, the landscape section of the <code>nameserver.ini</code> is also changed.</p>

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

[Configuring SAP HANA System Properties \(INI Files\) \[page 291\]](#)

[Database-Specific Configuration Parameters \[page 293\]](#)

10.2.4.1.8 Naming Conventions for Backups

When you plan your backup strategy, you need to be familiar with the naming conventions and recommendations for file-based data backups, delta backups, and third-party backup tools.

Related Information

[Naming Conventions for Data Backups \[page 1266\]](#)

[Naming Conventions for Delta Backups \[page 1269\]](#)

[Naming Conventions for Log Backups \[page 1270\]](#)

[Naming Conventions for the Backup Catalog \[page 1271\]](#)

[Temporary Names for File-Based Backups \[page 1271\]](#)

10.2.4.1.8.1 Naming Conventions for Data Backups

When you plan your backup strategy, you need to be familiar with the naming conventions and recommendations for data backups.

Each data backup name is comprised of the following elements:

```
<path><prefix>_<suffix>
```

i Note

The naming conventions apply to data backups created in the file system and data backups created using third-party tools. With third-party tools, you cannot change the backup path.

Elements of Data Backup Names

Name Element	Description
<code><path></code> For example: <code></backup/data/></code>	<p>Optional.</p> <p>For file-based backups:</p> <p>If no complete path is specified, the default backup destination is prepended to the backup name.</p> <p>For backups created using third-party tools:</p> <p>A named pipe is created in the file system for the system database or tenant database:</p> <ul style="list-style-type: none">• <code>/usr/sap/<SID>/SYS/global/hdb/backint/SYSTEM</code>• <code>/usr/sap/<SID>/SYS/global/hdb/backint/DB_<tenant_database_name></code> <p>A third-party backup tool reads data to be backed up from the named pipe, and writes the backup data in accordance with the tool configuration.</p>

Name Element	Description
<prefix>	<p>Optional.</p> <p>You can use the prefix proposed by the system or you can specify a different prefix for the backup name.</p>
	<p>i Note</p> <p>To be able to more easily identify archived file-based backups, it is strongly recommended that you use a unique prefix for each data backup name. For example, a timestamp. By default, the name of each scheduled backup is prefixed with the timestamp of the start of the backup. The placeholders [date] and [time] are automatically converted to the current timestamp.</p> <p>If you use the same prefixes, it is recommended that you replicate a data backup to a new destination as soon as the backup is created. Otherwise, an existing complete data backup with the same name will be overwritten by the next data backup.</p> <p>For backups created using third-party tools, data backups are not overwritten. The Backint for SAP HANA interface is able to identify multiple versions of backups with the same name.</p> <p>Nevertheless, for easier identification and versioning, it is strongly recommended to assign unique backup names. .</p>
	<p>i Note</p> <p>It is not possible to change the prefix of a backup after it has been created.</p>
Suffix	<p>To each backup name, the system adds a suffix that indicates the volume ID and the partition ID.</p> <p>As this is done for each service that is included in the backup, you only need to specify the name (<path><prefix>) for all the backups on all the hosts in the system. The next time a service is backed up, the system assigns the same suffix to the backup to that service.</p>

i Note

Once backups have been created, it is strongly recommended that you **do not change** their names.

When backups are created, their names are stored in the backup catalog. For a recovery, specific backup components are located using the names stored in the backup catalog. If the name of a backup was changed after it was recorded in the backup catalog, it will not be possible to locate it using the backup catalog, and it will not be possible to use it for a recovery.

You can copy or move file-based backups to a different location. If you use a moved backup for a recovery, you need to specify its current location.

❖ Example

Names for Parts of a Data Backup

During backup, each service backs up its data to the specified backup destination.

Below is an example of a set of backups from one data backup.

```
</backup/data/COMPLETE_DATA_BACKUP_databackup_0_1>
```

```
</backup/data/COMPLETE_DATA_BACKUP_databackup_1_1>
```

```
</backup/data/COMPLETE_DATA_BACKUP_databackup_2_1>
```

```
<...>
```

In the above example, the `<path>` is `</backup/data/>`, the `<prefix>` is `<COMPLETE_DATA_BACKUP>`. `<databackup_0_1>` is the suffix, which is automatically added by the system. In the suffix, `<0>` is the volume ID, and `<1>` is the partition ID

10.2.4.1.8.2 Naming Conventions for Delta Backups

When you plan your backup strategy, you need to be familiar with the naming conventions and recommendations for differential and incremental backups.

Structure of File Names for Delta Backups

	Differential Backups	Incremental Backups
Prefix:	User-defined. A timestamp is recommended. For example: 2018-01-23	User-defined. A timestamp is recommended. For example: 2018-01-23
String:	databackup_differential	databackup_incremental
Backup ID:	The backup ID of the full data backup that the differential backup is based on	The backup ID of the full data backup or the delta backup that the incremental backup is based on
Delta Backup ID:	ID of the differential backup	ID of the incremental backup
Volume ID:	Volume ID as with complete data backups	Volume ID as with complete data backups
Partition ID:	Partition ID as with complete data backups	Partition ID as with complete data backups

❖ Example

File Names for Differential Backups

The SQL statement `BACKUP DATA DIFFERENTIAL USING FILE ('2018-01-23')` creates a differential backup based on the previously created full data backup.

Example names of incremental backup files:

```
2018-01-23_databackup_differential_1426237023821_1426237780534_0_1
```

```
2018-01-23_databackup_differential_1426237023821_1426237780534_1_1
```

```
2018-01-23_databackup_differential_1426237023821_1426237780534_2_1
```

```
2018-01-23_databackup_differential_1426237023821_1426237780534_3_1
```

❖ Example

File Names for Incremental Backups

The SQL statement `BACKUP DATA INCREMENTAL USING FILE ('2018-01-23')` creates an incremental backup based on the previously created full data backup or differential backup.

Example names of incremental backup files:

```
2018-01-23_databackup_incremental_1426237023821_1426237028496_0_1
```

```
2018-01-23_databackup_incremental_1426237023821_1426237028496_1_1
```

```
2018-01-23_databackup_incremental_1426237023821_1426237028496_2_1
```

```
2018-01-23_databackup_incremental_1426237023821_1426237028496_3_1
```

10.2.4.1.8.3 Naming Conventions for Log Backups

Log backup names are generated automatically. Unlike the data backup names, no parts of the log backup names are user-defined.

The names of log backups are assigned in accordance with specific naming conventions.

Each log backup name comprises the following elements:

```
<log_backup>__<volume ID>_<log partition ID>_<first redo log position>_<last redo log position>.<backup_ID>
```

The elements of log backup names are separated by an underscore. A period ('.') separates the appended backup ID from the log name.

Elements of Log Backup Names

Name Element	Description
<log_backup>	All log backups begin with the string <log_backup>.

Name Element	Description
<code><volume ID></code>	The volume ID for the SAP HANA service. For example, name server, index server, script server, or XS engine.
<code><log partition ID></code>	Only one log partition is supported for each service.
<code><first redo log position></code>	The oldest entry in log backup
<code><last redo log position></code>	The most recent entry in the log backup
<code><backup_ID></code>	Uniquely identifies the log backup

i Note

`<backup_ID>` is calculated automatically by SAP HANA. `<backup_ID>` is only used for file-based backups, not for backups with third-party backup tools.

❖ Example

A log backup name could look like this:

```
<log_backup_1_0_1234567_1238567.1380740407446>
```

10.2.4.1.8.4 Naming Conventions for the Backup Catalog

Different names are assigned to the backup catalog with file-based backups and when using a third-party backup tool.

The backup catalog is assigned a name in the following format:

```
log_backup_0_0_0_0.<Backup_ID>
```

With a third-party tool, the backup catalog is assigned a name in the following format:

```
log_backup_0_0_0_0
```

10.2.4.1.8.5 Temporary Names for File-Based Backups

When file-based backups are written, they are first written using a temporary name.

After a part of a backup has been written successfully, it is renamed to the final name used for the SAP HANA service. Existing backups with the same name are only overwritten after the backup for the service was completed successfully.

i Note

If existing backups are overwritten by backups with the same names, at least **twice the space** in the backup location is needed, because the old backup and the new backup exist for a time in parallel.

10.2.4.1.9 Diagnosis Files for Backup and Recovery

The `backup.log` and `backint.log` files record information about backups. This information can be used to diagnose errors.

→ Tip

As more data is written to `backup.log` and `backint.log`, the files grow, but their increased size does not impact database performance. If `backup.log` or `backint.log` do become too big for the available disk space, you can safely delete either file as required.

Related Information

[backup.log \[page 1272\]](#)

[backint.log \[page 1273\]](#)

10.2.4.1.9.1 backup.log

The `backup.log` file records information about data backups, log backups, and the backup catalog.

`backup.log` also records information about recovery operations.

i Note

The SQL statement used for a recovery is recorded in `backup.log`. For a point-in-time recovery, the point in time is specified in the SQL statement as **UTC**.

The time at which the recovery was started and completed is recorded in `backup.log` as **local server time**, not UTC.

i Note

For a point-in-time recovery, the point in time specified in the SQL statement may be different from the point in time that was actually reached in the recovery. This is because the point in time that was actually reached in the recovery is that of the most recent global COMMIT to the database that was recovered.

For example, a point in time of **13:15** was specified for a recovery. SAP HANA interprets this time as UTC. The point in time reached could be **13:28:56+01:00** (local server time), which would be **12:28:56** as UTC.

The point in time recorded for a point-in-time recovery is the same regardless of whether the backups are from the file system or from a third-party backup tool.

Display the Content of backup.log

You can display the content of `backup.log` as follows:

SAP HANA cockpit

Choose [▶ View trace and diagnostic files](#) [▶ <system database>](#) [▶ Database Diagnostic Files](#) [▶ <host>](#) [▶ other](#) [▶](#).

SAP HANA studio

In the Backup Console on the *Overview* tab, choose [Open Log File](#).

In the Administration editor, go to the *Diagnosis Files* tab.

10.2.4.1.9.2 backint.log

`backint.log` contains information about the activities of the `Backint` agent.

The `Backint` agent is part of a third-party backup tool.

`backint.log` records all the parameters used to call the `Backint` agent, and the values returned. Each time the `Backint` agent is called, the command parameters and the return code are appended to `backint.log`.

`backint.log` includes the content of the following files:

- `Backint` input file
This file is created when the `Backint` agent is started.
- `Backint` output file
The `Backint` agent writes its output to this file.

The contents of the command file and the output file are copied to `backint.log`.

Display the Content of backint.log

You can display the content of `backint.log` as follows:

SAP HANA cockpit

Choose [▶ View trace and diagnostic files](#) [▶ <system database>](#) [▶ Database Diagnostic Files](#) [▶ <host>](#) [▶ other](#) [▶](#).

10.2.4.2 Configure Backups

Using SAP HANA cockpit, you can display an overview of the active backup and recovery configuration settings for each database, and change the default backup and recovery configuration settings for all the tenant databases and the system database.

Prerequisites

You have the authorization DATABASE ADMIN.

For more information, see *Authorizations for Backup and Recovery*.

Changes to configuration settings are made through the system database.

Procedure

1. From SAP HANA cockpit, select a system database, then go to *Overall Tenant Statuses* and open the overview of databases.

An overview of the databases is displayed.

2. Choose *Configure Backup*.

An overview of the current systemwide backup parameter settings is displayed.

The log mode and the backup encryption status is displayed for each database.

For more information about SAP HANA encryption, see *SAP HANA Backup Encryption*.

To navigate to a specific settings group, you can choose *Backint Response Timeout*, *Catalog Settings*, *Log Settings*, *Data Backup Settings*, or *Restrictions for Tenant Database Users*.

3. To change configuration settings, choose *Edit*.

i Note

When you edit a parameter group, the Backup Configuration app checks whether the parameter settings are within the recommended range. If a parameter has been changed to a non-recommended setting, a setting within the recommended range is proposed. You have the option to save or discard the recommended parameter setting.

To reset a parameter to the default setting, choose *Reset to Default*.

To retain a change, choose *Save*.

When you save, the change takes effect immediately.

For more information about the configuration options, see *Backup Configuration Settings*.

Related Information

[Backup Configuration Settings \[page 1275\]](#)

[Authorizations for Backup and Recovery \[page 1244\]](#)

[Log Modes \[page 1287\]](#)

[SAP HANA Backup Encryption \[page 1252\]](#)

10.2.4.2.1 Backup Configuration Settings

Using SAP HANA cockpit, you can display and change the default backup configuration settings for the system database and the tenant databases.

The backup configuration settings are described in the following sections.

Backint Settings

The parameter options for third-party backup tools are only available if the `Backint` agent is installed.

Backint Parameter Files

If required by the third-party backup tool, you can specify `Backint` parameter files for data backup and for log backups. The content and syntax of the parameter files is tool-specific and defined by the tool vendor.

For more information, see the vendor documentation for the third-party backup tool.

i Note

If you disable `Backint`, check that the destination used for file-based backups is correct.

Setting	Description
<i>Use the Same Parameter File for All</i>	You can use the same <code>Backint</code> parameter file for data backups, log backups, and for backups of the backup catalog.
<i>Data Backup, Log Backup, Catalog Backup</i>	You can specify a different <code>Backint</code> parameter file for data backups, log backups, and for backups of the backup catalog.

i Note

To use a parameter file, there needs to be a symbolic link pointing from `/usr/sap/<SID>/SYS/global/hdb/opt/hdbconfig/` to the actual parameter file in the directory.

If a new host is added, ensure that the database services have access to the Backint agent and the parameter file.

i Note

To specify parameter files for the databases in a high isolation system, you need to work from the system database.

With **high isolation**, the settings are configured separately for the system database and the tenant databases. For each database, you can use a different Backint parameter file for data backups, log backups, and for backups of the backup catalog.

In an SAP HANA database with many tenant databases, many Backint parameter files may be needed to ensure high isolation.

For more information, see *Isolation Level High for Backups and Third-Party Backup Tools*.

Backint Response Timeout

Setting	Description
<i>Set Timeout (Catalog and Log Backups)</i>	<p>Enable a timeout for the connection to the third-party backup tool. The timeout is measured from the time of the first request to the Backint for SAP HANA agent.</p> <div data-bbox="820 1084 1401 1189"><p>i Note</p><p>This timeout is reset when data is transferred.</p></div>
<i>Timeout After (Minutes)</i>	<p>Specify the timeout.</p> <p>The default timeout is 10 minutes. If you do not specify a timeout, the default value is used.</p> <p>You can change the setting in increments of 5 minutes up to a maximum of 30 minutes.</p> <p>For more information, see <i>Timeout for Log Backups (Backint)</i>.</p> <p>If the Backint process terminates as the result of a timeout, it may be recorded in the <code>backint.log</code> as having terminated with an error.</p> <p>For more information about <code>backint.log</code>, see <i>backint.log</i>.</p>

Backup Catalog Settings

Setting	Description
<i>Destination Type</i>	The destination type can be the file system or a third-party backup tool.
<i>Location</i>	<p>For Backint, the backup data is written through the third-party backup tool. You cannot change the location.</p> <p>For file system backups, by default, the backup catalog is backed up to the same location as the log backups: \$(DIR_INSTANCE) /backup/log</p> <p>You can specify an alternative location to which to write backups of the backup catalog.</p> <p>For more information, see <i>Destination for Backups of the Backup Catalog</i>.</p>

Log Settings

Log Mode

Setting	Description
<i>Log Mode</i>	<p>SAP HANA uses two log modes: <code>normal</code> and <code>overwrite</code>.</p> <p>By default, SAP HANA runs in log mode <code>normal</code>.</p> <div style="border: 1px solid #0070c0; padding: 5px;"><p>→ Tip</p><p>If you change the log mode from <code>overwrite</code> – where log backups are not written – to log mode <code>normal</code>, you must create a full data backup to ensure that log backups are written again, and that the database can be recovered to the most recent point in time.</p></div> <p>For more information, see <i>Change Log Modes</i>.</p>

Log Backup

Setting	Description
Create Log Backups	<p>Enable or disable automatic log backups.</p> <p>If Create Log Backups is disabled, the other settings in the parameter group cannot be changed.</p> <p>To enable log backups, the log mode must be set to <code>normal</code>. In log mode <code>overwrite</code>, you cannot change the log backup settings.</p> <p>For more information, see Log Modes.</p> <div style="border: 1px solid orange; padding: 5px;"><p>⚠ Caution</p><p>During normal system operation (log mode <code>normal</code>), it is strongly recommended that you enable automatic log backups. When log segments are backed up, the space they occupied in the log area can be freed. SAP HANA can overwrite the newly freed space in the log area with new log entries. In this way, automatic log backups can prevent the log area from filling. If automatic log backups are disabled, the log area grows until the file system is full. If the file system is full, and no more log segments can be created, the database freezes.</p></div>
Use Consolidated Backups	<p>To improve the performance of log backups, SAP HANA can write multiple log segments of a service to a single consolidated log backup.</p> <p>If you do not use consolidated log backups, each log segment is backed up to its own backup.</p> <p>For more information, see Consolidated Log Backups.</p>
Maximum Size (GB)	<p>You can configure the maximum size of consolidated log backups in increments of 8 GB.</p> <p>The default value is 16 GB.</p> <p>This means that one backup operation creates consolidated log backups with a maximum size of 16 GB.</p> <p>The minimum size is 8 GB. The maximum size allowed is 64 GB.</p> <p>For more information, see Consolidated Log Backups.</p>
Destination Type	<p>The destination type can be the file system or a third-party backup tool.</p>

Setting	Description
<i>Location</i>	<p>For Backint, the backup data is written through a third-party backup tool. You cannot change the location.</p> <p>For file system backups, by default, SAP HANA log segments in the log area are backed up to: <code>\$(DIR_INSTANCE) / backup/log</code></p> <p>You can specify an alternative location.</p>
<i>Back Up Logs</i>	<p>Specify when to back up the log segments.</p> <p>At latest after specified time limit: If log segments become full, they are backed up immediately, even if the log backup time limit has not been reached.</p> <p>(This is the same as log backup interval mode: immediate)</p> <p>Only after specified time limit: Log segments are backed up after the time limit you specify. This means that log segments are not automatically backed up if the log segments become full.</p> <p>(This is the same as log backup interval mode: service)</p> <p>For more information, see <i>Set the Interval Mode for Log Backups</i>.</p>
<i>Time Limit (Minutes)</i>	<p>Specify the time limit for log backups.</p> <p>By default, the time limit is 15 minutes.</p> <p>The maximum permitted time limit is 1440 minutes (24 hours).</p>

Data Backup Settings

Data Backup

Setting	Description
<i>Destination Type</i>	The destination type can be the file system or a third-party backup tool.

Setting	Description
<i>Location</i>	<p>Backint: The backup data is written through a third-party backup tool. You cannot change the location.</p> <p>File system backups: By default, the backups are written to the following location: <code>\$(DIR_INSTANCE) /backup/data</code></p> <p>You can specify an alternative location to which to write data backups.</p> <p>For more information, see <i>Parameters for Data Backup Settings</i>.</p>
<i>Limit Maximum Size (File System Backups)</i>	<p>For file system backups, you can specify a maximum file size.</p> <p>The maximum file size applies to the data backups of all services.</p> <p>If the size of a data backup file for a service exceeds the specified limit, SAP HANA splits the file into multiple smaller files.</p>
<i>Maximum Size (GB)</i>	<p>You can specify the maximum file size of data backups in increments of 50 GB up to 2000 GB.</p> <p>The actual size of data backups may be smaller than the specified maximum size.</p>

Setting	Description
Parallel Streams (Backint Backups)	<p>When creating a data backup, a third-party backup tool can use multiple channels in parallel to write the backup data for each service.</p> <p>By default, SAP HANA uses one channel for data backups. If required, you can configure SAP HANA to use additional channels. When multiple channels are used, SAP HANA distributes the data equally across the available channels.</p> <p>You can use up to 32 parallel streams for backups using a third-party tool.</p> <p>All the parts of a multistreamed backup are approximately the same size.</p> <p>For more information, see <i>Multistreaming Data Backups with Third-Party Backup Tools</i>.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>To create multistreamed data backups, the third-party backup tool must also be configured to use multiple channels with good performance.</p> <p>For more information about the configuration of the backup tool, consult the vendor documentation.</p> </div>

Data Backup Scheduler

Setting	Description
Enable Data Backup Scheduler	<p>To be able to execute scheduled data backups, the data backup scheduler must be enabled. (If the scheduler is not enabled, you can still schedule backups, but they cannot be executed.)</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>A backup of a tenant database must be scheduled through the tenant database itself. A backup of a tenant database cannot be scheduled through the system database.</p> </div> <p>For more information, see <i>Schedule Backups</i>.</p>
Assigned Database User	<p>The user that activates the backup schedules is also used to execute the backups. If the data backup scheduler is running, the associated user is displayed.</p>

Restrictions for Tenant Database Users

Setting	Description
<i>Users Can Create Backups</i>	<p>You can prevent all users of a tenant database from creating backups.</p> <p>For more information about the authorizations needed, see <i>Authorizations for Backup and Recovery</i>.</p>
<i>Users Have Free Choice of Backup Destination</i>	<p>i Note</p> <p>For file-system backups only.</p> <p>For users who are allowed to create backups, you can allow backups to be created:</p> <ul style="list-style-type: none">• In any directory• Only in the default backup destination or a subpath of the default backup destination. <p>Users can create new subdirectories for backups below the default backup destination.</p> <p>i Note</p> <p>After you change the permission for the backup destination for all tenant databases, the database must be re-started.</p>

Related Information

[Working with Third-Party Backup Tools \[page 1303\]](#)

[Isolation Level High for Backups and Third-Party Backup Tools \[page 1308\]](#)

[Timeout for Log Backups \(Backint\) \[page 1307\]](#)

[Multistreaming Data Backups with Third-Party Backup Tools \[page 1306\]](#)

[Delta Backups and Third-Party Backup Tools \[page 1311\]](#)

[Parameters for Backing Up the Backup Catalog \[page 1297\]](#)

[Destination for Backups of the Backup Catalog \[page 1299\]](#)

[Accumulated Backups of the Backup Catalog \[page 1301\]](#)

[Log Modes \[page 1287\]](#)

[Change Log Modes \[page 1290\]](#)

[Consolidated Log Backups \[page 1296\]](#)

[Set the Interval Mode for Log Backups \[page 1291\]](#)

[Parameters for Data Backup Settings \[page 1283\]](#)

[Schedule Backups \[page 1325\]](#)

[Authorizations for Backup and Recovery \[page 1244\]](#)

[backint.log \[page 1273\]](#)

10.2.4.2.2 Backup Configuration Parameters

In addition to displaying and changing backup configuration settings using SAP HANA cockpit and SAP HANA studio, you can also configure the underlying parameters.

For more information about changing parameter values, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

[Parameters for Data Backup Settings \[page 1283\]](#)

[Parameters for Log Backup Settings \[page 1286\]](#)

[Parameters for Backing Up the Backup Catalog \[page 1297\]](#)

[Configure a Third-Party Backup Tool \[page 1304\]](#)

[Backing Up Customer-Specific Configuration Settings \[page 1264\]](#)

10.2.4.2.2.1 Parameters for Data Backup Settings

Data backups can be written to different destinations. However, all the parts of the same data backup are written to the same destination.

For file-based data backups, you can change the default destination.

For backups created using third-party backup tools, the default backup destination cannot be changed.

Related Information

[Change the Default Destination for File-Based Data Backups \[page 1284\]](#)

[Naming Conventions for Backups \[page 1266\]](#)

[Configure a Third-Party Backup Tool \[page 1304\]](#)

[Configure Backups \[page 1274\]](#)

10.2.4.2.2.1.1 Change the Default Destination for File-Based Data Backups

Each time you start a file-based data backup, you have the option to change the default backup destination or to specify a different destination for the current backup only.

Context

By default, file-based data backups are written to `$DIR_INSTANCE/backup/data`.

→ Tip

For file-based backups, it is recommended that you create the destination directory structures **before the backup is started**.

Procedure

1. Locate the parameter `basepath_databackup` in the `persistence` section of the parameter file.

For a tenant database, the parameter file is `global.ini`.

For a system database, the parameter file is `nameserver.ini`.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

→ Tip

For improved data safety, it is recommended that you specify a path to a secure backup destination. The data area, log area, and backups should never be on the same physical storage devices.

2. Open the change dialog.
3. Specify the new default destination.
4. Save.

Results

The change takes effect immediately.

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

10.2.4.2.2.1.2 Set the Maximum File Size for File-Based Backups

For file-based data backups, you may need to limit the maximum size of a single backup file. For example, due to file system limitations.

Context

If the size of a data backup file for a service exceeds the specified limit, SAP HANA splits the file into multiple smaller files.

Procedure

1. Locate the parameter `data_backup_max_chunk_size` in the `backup` section of the parameter file.

For a tenant database, the parameter file is `global.ini`.

For a system database, the parameter file is `nameserver.ini`.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

2. Open the change dialog.
3. Specify the new maximum size.

You can specify the maximum file size of data backups up to 2000 GB.

The actual size of data backups may be smaller than the specified maximum size.

i Note

If existing backups are overwritten by backups with the same names, at least **twice the space** in the backup location is needed, because the old backup and the new backup exist for a time in parallel.

4. Save.

i Note

Changes take effect immediately.

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

10.2.4.2.2.2 Parameters for Log Backup Settings

You can configure parameters to control log backups.

Related Information

[Enable and Disable Automatic Log Backup \[page 1286\]](#)

[Log Modes \[page 1287\]](#)

[Change Log Modes \[page 1290\]](#)

[Set the Interval Mode for Log Backups \[page 1291\]](#)

[Change the Log Backup Interval \[page 1292\]](#)

[Change the Log Backup Destination Type \[page 1294\]](#)

[Change the Log Backup Destination \[page 1295\]](#)

[Consolidated Log Backups \[page 1296\]](#)

[Configure Backups \[page 1274\]](#)

10.2.4.2.2.2.1 Enable and Disable Automatic Log Backup

By default, SAP HANA creates redo log backups automatically at regular intervals. You can disable and enable automatic log backups.

Prerequisites

To enable automatic log backups, the log mode must be set to `normal`.

For more information, see *Log Modes*.

⚠ Caution

During normal system operation (log mode `normal`), it is strongly recommended that you enable automatic log backups. When log segments are backed up, the space they occupied in the log area can be freed. SAP HANA can overwrite the newly freed space in the log area with new log entries. In this way, automatic log backups can prevent the log area from filling. If automatic log backups are disabled, the log

area grows until the file system is full. If the file system is full, and no more log segments can be created, the database freezes.

Procedure

1. Locate the parameter `enable_auto_log_backup` in the `persistence` section of the parameter file.

For a tenant database, the parameter file is `global.ini`.

For a system database, the parameter file is `nameserver.ini`.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

2. Open the change dialog.
3. Specify whether to enable or disable automatic log backups.

The default setting is `yes` (automatic log backup is active).

You can specify either `yes` or `no` to enable or disable automatic log backups.

To reset to enable automatic log backups, choose *Restore Default*.

4. Save.

Results

The change takes effect immediately.

If any log backups are running, they will first be completed before automatic log backups are disabled.

Related Information

[Log Modes \[page 1287\]](#)

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

[Configure Backups \[page 1274\]](#)

10.2.4.2.2.2 Log Modes

SAP HANA can run either in log mode `normal` or `overwrite`.

After installation, SAP HANA temporarily runs in log mode `overwrite`.

In log mode `overwrite`, no log backups are created.

Log mode `overwrite` ensures that the log area does not grow excessively.

After you create the first full data backup, SAP HANA automatically switches to the default log mode `normal`.

→ Tip

If you change the log mode from `overwrite` – where log backups are not written – to log mode `normal`, **you must create a full data backup** to ensure that log backups are written again, and that the database can be recovered to the most recent point in time.

SAP HANA Log Modes

Log Mode	Description
<code>normal</code> (Default)	<p>In log mode <code>normal</code>, log segments are backed up automatically if automatic log backups are enabled.</p> <p>For more information, see <i>Enable and Disable Automatic Log Backup</i> and <i>Changing the Backup Configuration Settings</i>.</p> <p>→ Tip</p> <p>Log mode <code>normal</code> is recommended to provide support for point-in-time recovery.</p> <p>After a log segment has been backed up, closed, and a savepoint has been written, the space it occupied can be freed. SAP HANA can then overwrite the newly freed space in the log area with new log entries. In this way, automatic log backups can prevent the log area from filling.</p> <p>⚠ Caution</p> <p>If the log area becomes full and no more log segments can be created in the file system, the database freezes. No more log entries can be written until a log backup has been completed, and the log segments are no longer needed to restart the database.</p>

Log Mode	Description
overwrite	<p>No log backups are created. When savepoints are written, log segments are immediately freed to be overwritten by new log entries.</p> <p>When log mode <code>overwrite</code> is active, the Log Backup Settings in the Backup Console cannot be changed.</p> <p>Log mode <code>overwrite</code> can be useful for installations that do not need to be backed up or recovered. For example, for test installations.</p> <div data-bbox="616 611 1394 936" style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>⚠ Caution</p> <p>Log mode <code>overwrite</code> is not recommended for production systems.</p> <p>With log mode <code>overwrite</code>, a point-in-time recovery is not possible. Delta backups created in log mode <code>overwrite</code> cannot be used for a point-in-time recovery.</p> <p>Only the following recovery option can be selected: Recover the database to a specific data backup or storage snapshot</p> </div> <div data-bbox="616 956 1394 1090" style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>Even if no log backups are written, with each data backup, the backup catalog is still backed up.</p> </div>

⚠ Caution

Do not delete log segments at operating system level, as this makes the log area unusable. As a consequence, the database may stop working immediately, and it will not be possible to restart the database.

For more information about switching log modes, see [Backup Configuration Settings](#) or [Change Log Modes](#).

Related Information

- [Backup Configuration Settings \[page 1275\]](#)
- [Change Log Modes \[page 1290\]](#)
- [Enable and Disable Automatic Log Backup \[page 1286\]](#)
- [Change a System Property in SAP HANA Studio \[page 301\]](#)
- [Savepoints and Redo Logs \[page 1230\]](#)
- [SAP Note 2375865 !\[\]\(5c31e34db2feb955590990a7829cf6bf_img.jpg\)](#)
- [Change the Log Backup Interval \[page 1292\]](#)
- [Set the Interval Mode for Log Backups \[page 1291\]](#)

10.2.4.2.2.2.1 Change Log Modes

You can switch between the SAP HANA log modes `normal` and `overwrite`.

Procedure

→ Tip

If you change the log mode from `overwrite` – where log backups are not written – to log mode `normal`, **you must create a full data backup** to ensure that log backups are written again, and that the database can be recovered to the most recent point in time.

1. Locate the parameter `log_mode` in the `persistence` section of the parameter file.

For a tenant database, the parameter file is `global.ini`.

For a system database, the parameter file is `nameserver.ini`.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

2. Open the change dialog.
3. Specify the new log mode.

The log mode can be either `normal` or `overwrite`.

For more information, see *Log Modes*.

To reset SAP HANA to the default log mode `normal`, choose *Restore Default*.

4. Save.

The change takes effect immediately.

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

[Log Modes \[page 1287\]](#)

10.2.4.2.2.3 Set the Interval Mode for Log Backups

You can specify an interval mode for log backups.

Context

→ Tip

You can also specify the log backup interval mode using the Backup Configuration App in SAP HANA cockpit or the Backup Console in SAP HANA studio.

Procedure

1. Locate the parameter `log_backup_interval_mode` in the `backup` section of the parameter file.

For a tenant database, the parameter file is `global.ini`.

For a system database, the parameter file is `nameserver.ini`.

For more information about how to change parameter settings, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

`log_backup_interval_mode` controls the use of the log backup timeout.

For more information, see *Change the Log Backup Interval*.

2. Specify the desired log backup interval mode.

You can set the following interval modes:

Log Backup Interval Mode	Description
<i>immediate</i> (Default)	<p>A log backup is created immediately after a log segment becomes full, or when the service-specific timeout for a log segment has been reached.</p> <p>If you have configured consolidated log backups (parameter <code>max_log_backup_size</code>), the log backup process does not wait for the configured size of log backups to be queued. If fewer log backups are queued, all the queued log backups are processed.</p> <p>For more information, see <i>Consolidated Log Backups</i>.</p>

Log Backup Interval Mode	Description
<i>service</i>	A log backup is created only after a service-specific timeout has been reached. This backup includes all the log segments in states writing , closed , and truncated .

3. *Save*.

Changes take effect immediately.

i Note

The log interval mode and timeout are included in the system view **M_BACKUP_CONFIGURATION**.

For more information, see *M_BACKUP_CONFIGURATION*.

Related Information

[Backup Configuration Settings \[page 1275\]](#)

[Change the Log Backup Interval \[page 1292\]](#)

[Timeout for Log Backups \(Backint\) \[page 1307\]](#)

[Consolidated Log Backups \[page 1296\]](#)

10.2.4.2.2.2.4 Change the Log Backup Interval

You can change the interval at which log backups are created.

Prerequisites

The log backup interval takes effect only if **automatic log backups** are enabled.

To enable automatic log backups, the log mode must be `normal`.

For more information, see *Log Modes*.

Context

Specifying an appropriate interval for log backups enables you to recover an SAP HANA database with a good Recovery Point Objective (RPO). In the event of database failure, the RPO is the maximum time span from

which data will be lost if the log area cannot be used for recovery, and if only data backups, delta backups, and log backups are available.

i Note

If the log segments become full before the log backup interval, the logs are backed up automatically.

→ Tip

You can specify the log backup interval using the Backup Configuration App in SAP HANA cockpit or the Backup Console in SAP HANA studio.

Procedure

Change the log backup interval.

1. Locate the parameter `log_backup_timeout_s` in the `backup` section of the parameter file.

For a tenant database, the parameter file is `global.ini`.

For a system database, the parameter file is `nameserver.ini`.

For more information about how to change parameter settings, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

2. Specify the desired log backup interval.

By default, the log backup interval is 15 minutes (900 s).

A log backup interval of 15 minutes (or less) is recommended for production systems.

For test systems, you can set a longer log backup interval, depending on what data loss is acceptable to you if a fault occurs.

i Note

If you specify a timeout of 0, log backups are created only when a log segment is full and when services are restarted.

3. Save.

i Note

If log segments become full, they are backed up immediately, even if the log backup interval has not been reached.

Results

The change takes effect immediately.

Related Information

[Log Modes \[page 1287\]](#)

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

[Set the Interval Mode for Log Backups \[page 1291\]](#)

10.2.4.2.2.2.5 Change the Log Backup Destination Type

You can change parameters to the destination type for file-based log backups.

Procedure

1. Locate the parameter `log_backup_using_backint` in the `backup` section of the parameter file.

For a tenant database, the parameter file is `global.ini`.

For a system database, the parameter file is `nameserver.ini`.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

2. Open the change dialog.
3. Specify the destination type.

Possible values:

- `true` (back up using Backint)
- `false` (back up to the file system)

By default, log backups are written to the file system.

i Note

The destination type *Backint* is only available if the Backint agent is installed.

4. Save.

The change takes effect immediately.

Related Information

[Destination for Backups of the Backup Catalog \[page 1299\]](#)

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

10.2.4.2.2.2.6 Change the Log Backup Destination

You can change parameters to the default destination for file-based log backups.

Context

i Note

The log backup destination can only be changed for file-based backups.

It is not possible to change the backup destinations for third-party backup tools.

Backups created using third-party backup tools always use the destination: `/usr/sap/<SID>/SYS/global/hdb/backupint`

A third-party backup tool reads data to be backed up from named pipes, and writes the data in accordance with the tool configuration. For a third-party backup tool, the only objects created in the file system are named pipes. Named pipes occupy no space in the file system.

Procedure

1. Locate the parameter `basepath_logbackup` in the `persistence` section of the parameter file.

For a tenant database, the parameter file is `global.ini`.

For a system database, the parameter file is `nameserver.ini`.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

2. Open the change dialog.
3. Specify the destination for log backups.

By default, file-based log backups are written to: `$DIR_INSTANCE/backup/log`

→ Tip

For improved data safety, it is recommended that you specify a path to a secure backup destination.

The data area, log area, data backups, and log backups should never be on the same physical storage devices.

4. Save.

Results

The change takes effect immediately.

10.2.4.2.2.2.7 Consolidated Log Backups

To improve the performance of log backups, SAP HANA can write multiple log segments of a service to a single consolidated log backup. You can configure the maximum size of consolidated log backups.

Context

Log segments in the log area are only released after a log backup was completed successfully. After log segments have been released, they can then be overwritten with new log segments.

If a single log backup operation takes a long time, during that time, several other log segments may be queued for backup. During periods of high load, log segments may be closed and queued for backup faster than a single backup operation can be completed.

If log segments are waiting to be backed up, there can be a delay in releasing them. As a result of this delay, the log area may grow.

If the log segments cannot be backed up and released faster than the log area is growing, the log area could become full.

To remedy this issue, SAP HANA can write all the log segments of a service that are ready to be backed up at a particular time to a single consolidated log backup.

i Note

This option is supported for both file-based backups and third-party backup tools.

Procedure

Configure the maximum size of a consolidated log backup to be processed by a single backup operation.

1. Locate the parameter `max_log_backup_size` in the `backup` section of the parameter file.

For a tenant database, the parameter file is `global.ini`.

For a system database, the parameter file is `nameserver.ini`.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

2. Open the change dialog.
3. Select the database or the host.

You can configure the size of a consolidated log backup for one or more tenant databases and for one or more hosts.

4. Set the size in GB.

The default value is **16** GB.

This means that one backup operation creates consolidated log backups with a maximum size of 16 GB.

A consolidated log backup that contains many log segments can become quite large.

The maximum size allowed is 64 GB.

To reset to the default value, choose *Restore Default*.

i Note

The behavior of the log backup process depends on whether the parameter `log_backup_interval_mode` is *immediate* or *service*.

For more information, see *Set the Interval Mode for Log Backups*.

In the SAP HANA log area, the log segments that are part of a log backup are only released after the last of the queued log segments has been backed up and if they are no longer needed to restart the database.

5. Save.

Results

The backup catalog is written and backed up **once** for each log backup.

i Note

In the backup catalog, separate entries are maintained for the log segments in a log backup. However, if you subsequently want to remove log backups in a queued group, you can only remove all the log backups in the group together, and then only if none of them is still needed for recovery.

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

[Set the Interval Mode for Log Backups \[page 1291\]](#)

10.2.4.2.2.3 Parameters for Backing Up the Backup Catalog

Each time a backup of any type is created, the backup catalog is backed up and versioned. In this way, the latest version of the backup catalog always contains the entire backup history.

Even in situations such as when `log_mode = overwrite` is set, where log backups are not created, the backup catalog is still backed up.

If the backup catalog is backed up using a third-party tool, the tool also handles the versioning of the backup catalog.

You can manually change the following settings for the backup catalog:

- The destination type
The destination type can be the file system or a third-party backup tool.
For more information, see *Destination Type for Backups of the Backup Catalog*.
- Where the backup catalog is backed up
For more information, see *Destination for Backups of the Backup Catalog*.
- Whether backups of the backup catalog are accumulated
For more information, see *Accumulated Backups of the Backup Catalog*.

Related Information

[Destination Type for Backups of the Backup Catalog \[page 1299\]](#)

[Destination for Backups of the Backup Catalog \[page 1299\]](#)

[Accumulated Backups of the Backup Catalog \[page 1301\]](#)

[Configure Backups \[page 1274\]](#)

[Configure a Third-Party Backup Tool \[page 1304\]](#)

[Log Modes \[page 1287\]](#)

10.2.4.2.2.3.1 Destination Type for Backups of the Backup Catalog

For backups of the backup catalog, you can configure the backup destination type.

Parameter to Configure the Destination Type for Backups of the Backup Catalog

Task	Parameter
Destination Type:	<p>Write backups of the backup catalog to the file system or using a third-party backup tool.</p> <p><code>catalog_backup_using_backint</code></p> <p>in <code>global.ini/[backup]/</code></p> <p>Values:</p> <ul style="list-style-type: none">• true (back up using Backint)• false (back up to the file system) <p>Default value: false (Backups of the backup catalog are written to the file system.)</p>

For more information about changing parameter settings, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

10.2.4.2.2.3.2 Destination for Backups of the Backup Catalog

For backups of the backup catalog, you can configure the backup destination.

By default, the backup catalog is backed up to the same destination as the **log backups**.

The location of the **log backups** is configured separately from the location of backups of the backup catalog.

⚠ Caution

If you change the default destination for the log backups, the backup catalog is **not automatically backed up to the same location**.

For this reason, if you change the default destination for the log backups, you must also check that the backup catalog is backed up to the desired destination.

For more information about changing parameter settings, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

Parameters to Configure the Destination for Backups of the Backup Catalog

Task	Parameter
File-based Backups:	<p>Specify the directory to which to write backups of the backup catalog.</p> <p><code>basepath_catalogbackup</code></p> <p><code>in global.ini/[persistence]/</code></p> <p>Default value: empty (Backups of the backup catalog are written to the default directory for log backups: <code>\$DIR_INSTANCE/backup/log</code>)</p>
Backint:	<p>Specify a Backint parameter file for the backup catalog.</p> <p><code>catalog_backup_parameter_file</code></p> <p><code>in global.ini/[backup]/</code></p> <p>Default value: empty (No Backint parameter file is used.)</p>

Considerations for Recovery

When you recover SAP HANA, you are prompted to select a directory or destination type for the backup catalog, regardless of the location of the backup catalog that is currently configured. SAP HANA then only searches the specified directory (not the subdirectories) or Backint, and selects the most recent backup catalog available there.

To recover SAP HANA using native SQL, the location of the backup catalog in the file system can be specified in the syntax for `RECOVER DATABASE` and `RECOVER DATA` using the clause `USING CATALOG PATH ('<path>')` or `USING CATALOG BACKINT`. If no location is specified, the location specified in the `.ini` files is used.

For more information, see *RECOVER DATABASE Statement (Backup and Recovery)* and *RECOVER DATA Statement (Backup and Recovery)* in the *SAP HANA SQL and System Views Reference*.

Considerations for System Copy

When you recover SAP HANA to create a system copy, it is not necessary to move old log backups and old backups of the backup catalog. The location of the backup catalog is specified explicitly.

Considerations for an SAP HANA Upgrade

During an upgrade from SAP HANA 1.0 to SAP HANA 2.0, customer-specific configuration settings are automatically retained.

→ Tip

Before an upgrade, if the destination for log backups was changed, you should check that the backup catalog is being backed up to the desired destination. Check the parameters `catalog_backup_using_backint`, `catalog_backup_parameter_file`, and `basepath_catalogbackup`.

Consideration for a New SAP HANA Installation

When you install a new SAP HANA system, ensure that the log backups and the backup catalog are BOTH backed up to the desired destination.

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

[Consolidated Log Backups \[page 1296\]](#)

[Recovering an SAP HANA Database \[page 1347\]](#)

[Copying a Database Using Backup and Recovery \[page 1374\]](#)

[Change the Log Backup Destination Type \[page 1294\]](#)

[Timeout for Log Backups \(Backint\) \[page 1307\]](#)

10.2.4.2.2.3.3 Accumulated Backups of the Backup Catalog

Each time a backup is created, the operation is recorded in the backup catalog, which is itself then backed up. If many data and log backups are created within a short time, the backup catalog would need to be backed up just as frequently.

If many backups of the backup catalog are queued to run, the most recently completed backup of the backup catalog will not reflect the most recent database backups. If many backups are waiting to be processed, this

can cause increased backup times, as a backup is only completed when it has been recorded in the backup catalog and the backup catalog has been backed up.

To address this issue, SAP HANA can accumulate changes to the backup catalog, and back them up together in one operation.

A new backup of the backup catalog would then include all the changes to the backup catalog that were made since the last backup of the backup catalog. Accumulating multiple backups of the backup catalog in this way has the advantage that fewer backups of the backup catalog are created.

Accumulated backups of the backup catalog are supported for both file system backups and third-party backup tools.

10.2.4.2.2.3.3.1 Disable Writing Accumulated Backups of the Backup Catalog

By default, writing accumulated backups of the backup catalog is enabled. You can disable writing accumulated backups of the backup catalog.

Procedure

1. Locate the parameter `enable_accumulated_catalog_backup` in the `backup` section of `global.ini`.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

2. Open the change dialog.
3. Disable (or enable) writing accumulated backups of the backup catalog.

To back up the backup catalog after each log backup, set the parameter to **false**.

The default setting is **true** (multiple log backups are accumulated to one backup of the backup catalog).

To reset to the default behavior, choose *Restore Default*.

4. Save.

i Note

Changes take effect immediately.

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

10.2.4.3 Working with Third-Party Backup Tools

Third-party backup tools can be fully integrated with SAP HANA to enable you to perform backup and recovery operations from SAP HANA cockpit, SAP HANA studio, and using native SQL.

Backint for SAP HANA Interface

In addition to the file system, you can back up and recover an SAP HANA database using an SAP-certified third-party tool that supports the `Backint for SAP HANA interface`, which is used to communicate with an SAP HANA database.

Each active host in a distributed SAP HANA system may have one or more volumes to be backed up. When `Backint for SAP HANA` is used to back up a database, several communication processes are started, one for each volume. Backint-based data backups and log backups can be created in parallel.

A third-party backup tool reads data to be backed up from named pipes, and writes the data in accordance with the tool configuration. For a third-party backup tool, the only objects created in the file system are named pipes. Named pipes occupy no space in the file system.

Prerequisites for Using Third-Party Backup Tools

- The implementation of the API of a third-party backup tool that uses the `Backint for SAP HANA` interface must be certified by SAP.
- You have a support contract with the tool vendor that permits you to use the third-party backup tool with SAP HANA.

More information:

- [SAP Note 1730932 \(Using Backup Tools with Backint for SAP HANA\)](#)
- For a current overview of certified third-party backup tools, go to the [SAP Certified Solutions Directory](#). Use the search term `HANA-BRINT` and select a partner name to display more details.
- [SAP Note 1730998](#) contains a list of backup tool versions that should **not** be installed or activated in an SAP HANA appliance.
- For more information about installing and configuring a third-party backup tool, consult the documentation provided by the tool vendor.

Related Information

[SAP Note 1730932](#) 

[SAP Note 1730998](#) 

[SAP Certified Solutions Directory](#) 

10.2.4.3.1 Configuring a Third-Party Backup Tool

After a third-party backup tool has been installed, you can back up and recover an SAP HANA database without making any further changes to the default configuration. However, you have the option to change some of the tool configuration settings.

For third-party backup tools, the following directories are used for both data backups and log backups:

- System database: `/usr/sap/<SID>/SYS/global/hdb/backupint/SYSTEMDB`
- Tenant database: `/usr/sap/<SID>/SYS/global/hdb/backupint/DB_<tenant_database_name>`

Note

The content of the backups is not necessarily written to these directories. The backup tool decides where the content of the backups is written to.

A third-party backup tool reads data to be backed up from named pipes, and writes the data in accordance with the tool configuration. For a third-party backup tool, the only objects created in the file system are named pipes. Named pipes occupy no space in the file system.

For third-party backup tools, the backup destinations cannot be changed. During a recovery, SAP HANA queries information about the backup destinations from the third-party backup tool.

Related Information

[Configure a Third-Party Backup Tool \[page 1304\]](#)

[Configure Backups \[page 1274\]](#)

10.2.4.3.1.1 Configure a Third-Party Backup Tool

SAP HANA offers several options for changing the configuration settings of a third-party backup tool.

Prerequisites

The Backint agent has been installed.

If the Backint agent is not installed, you cannot change the Backint parameter files.

SAP HANA expects the Backint agent executable (`hdbbackint`) to be in the following path:

`/usr/sap/<SID>/SYS/global/hdb/opt/hdbbackint`

If the Backint agent executable is not installed in this path, a symbolic link must be created during the installation of a third-party backup tool. This symbolic link points from `/usr/sap/<SID>/SYS/global/hdb/opt/hdbbackint` to the actual location of the Backint agent executable.

i Note

You cannot change the `backint` agent using SAP HANA cockpit or SAP HANA studio.

Procedure

You can display and change configuration parameters using SAP HANA cockpit and SAP HANA studio.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

To display all the parameters for third-party backup tools, search for `backint` from the system properties overview.

The following parameters are available for third-party backup tools:

Option	Parameter
Specify a timeout for log backups for third-party tools.	<code>backint_response_timeout</code> For more information, see <i>Timeout for Log Backups (Backint)</i> .
Write backups of the backup catalog a third-party backup tool.	<code>catalog_backup_using_backint</code> For more information, see <i>Destination for Backups of the Backup Catalog</i> .
Enable or disable log backup using a third-party tool.	<code>log_backup_using_backint</code> For more information, see <i>Change the Log Backup Destination Type</i> .
Specify the number of channels to be used for multistreaming.	<code>parallel_data_backup_backint_channels</code> For more information, see <i>Multistreaming Data Backups with Third-Party Backup Tools</i> and <i>Prerequisites: Recovery Using Multistreamed Backups</i> .

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

[Configure Backups \[page 1274\]](#)

Parameters for Third-Party Backup Tools

[Timeout for Log Backups \(Backint\) \[page 1307\]](#)

[Destination for Backups of the Backup Catalog \[page 1299\]](#)

[Change the Log Backup Destination Type \[page 1294\]](#)

[Multistreaming Data Backups with Third-Party Backup Tools \[page 1306\]](#)

10.2.4.3.1.1.1 Multistreaming Data Backups with Third-Party Backup Tools

When creating a data backup, a third-party backup tool can use multiple channels to write the backup data for each service.

Context

This capability allows you to distribute backup data in parallel to multiple devices.

For more information, consult the documentation for your third-party backup tool.

By default, SAP HANA uses one channel for data backups. If required, you can configure SAP HANA to use additional channels. When multiple channels are used, SAP HANA distributes the data equally across the available channels. All the parts of a multistreamed backup are approximately the same size.

i Note

To create multistreamed data backups, the third-party backup tool must also be configured to use multiple channels with good performance.

For more information about the configuration of the backup tool, consult the vendor documentation.

Procedure

Specify the number of channels to be used for multistreaming.

1. Locate the parameter `parallel_data_backup_backint_channels` in the backup section of `global.ini`.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

2. Open the change dialog.
3. Specify the number of channels to be used for multistreaming.

The default value of `parallel_data_backup_backint_channels` is 1.

A value of 1 means that data backup with third-party backup tools is done through ONE channel.

The maximum number of channels permitted is 32 for each service.

i Note

The number of multistreaming channels applies to all data backup services larger than 128 GB. Data backup services smaller than 128 GB always use only one channel.

i Note

Each additional channel requires an IO buffer of 512 MB. Ensure that increasing the number of channels does not have a negative impact on memory consumption.

To change the buffer size, use parameter `data_backup_buffer_size`.

4. Save.

i Note

Changes take effect immediately.

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

[Prerequisites: Recovery Using Multistreamed Backups \[page 1335\]](#)

10.2.4.3.1.1.2 Timeout for Log Backups (Backint)

If the Backint agent does not respond for a user-specified time when writing log backups, SAP HANA cancels the Backint process.

The timeout for log backups for third-party tools is specified using the following parameter:

```
backint_response_timeout
```

The default timeout is 600 s.

i Note

This parameter is used for log backups and backups of the backup catalog.

If the Backint process terminates as the result of a timeout, it may be recorded in the `backint.log` as having terminated with an error.

The following error message is written to the trace file for the service:

```
Backint did not respond for 600 seconds, killing pid
```

Related Information

[Set the Interval Mode for Log Backups \[page 1291\]](#)

[SAP Note 2571163](#)

10.2.4.3.1.1.3 Isolation Level High for Backups and Third-Party Backup Tools

SAP HANA supports high isolation for third-party backup tools.

In an SAP HANA database, it is necessary to ensure that a tenant database cannot access the backups of other tenant databases. If your third-party backup tool does not support high isolation, you can set up separate Backint parameter files for each tenant database. These Backint parameter files are owned and managed by the operating system user (`<sid>adm`), which has read and write access. The tenant databases have only read access to the Backint parameter file through the tenant-specific group.

For more information about the access permissions required for system and tenant databases, see *Set the Isolation Level to High for Backups with Third-Party Backup Tools*.

i Note

To ensure high isolation in an SAP HANA database with many tenant databases, many Backint parameter files may be needed.

→ Tip

Check with your third-party backup tool vendor whether any tool-specific restrictions apply.

Related Information

[Set the Isolation Level to High for Backups with Third-Party Backup Tools \[page 1309\]](#)

[Increase the System Isolation Level \[page 202\]](#)

[Database Isolation \[page 204\]](#)

10.2.4.3.1.1.3.1 Set the Isolation Level to High for Backups with Third-Party Backup Tools

By default, an SAP HANA tenant database has isolation level low. You can increase the isolation level to high to ensure that one tenant database cannot access the backups of other tenant databases.

For the System Database

Procedure

1. On operating system level, create a database-specific directory for the Backint parameter files.
The database-specific directory must be owned by the operating system user `<sid>adm` and the group `sapsys`.
2. Assign the access permissions 700 to the directory.
700 allows user `<sid>adm` read, write, and execute access to the directory; the group has no access permission; others have no access permission.
3. In the database-specific directory, create a Backint parameter file for backups.
If required, also create a database-specific Backint parameter file for log backups.
The parameter file must be owned by the operating system user `<sid>adm` and group `sapsys`.
4. Assign the access permissions 600 to the parameter file.
600 allows user `<sid>adm` read and write access to the file; the group has no access permission; others have no access permission.
5. Use the parameters `data_backup_parameter_file`, `log_backup_parameter_file`, and `catalog_backup_parameter_file` to specify the Backint parameter files.

For Each Tenant Database

Context

To grant the system administrator access to the tenant database backup files and directories, you need to add the `<sid>adm` user to the operating system group of each tenant. For more information, see *File and Directory Permissions with High Isolation* in the *SAP HANA Administration Guide*.

Procedure

1. On operating system level, create a database-specific directory for the Backint parameter files.

The database-specific directory must be owned by the operating system user `<sid>adm` and by the group of the tenant database.

2. Assign the access permissions 750 to the directory.

750 allows user `<sid>adm` read, write, and execute access to the directory; the group has read and execute permission; others have no access permission.

3. In each tenant-specific directory, create a tenant-specific Backint parameter file for backups.

If required, also create a tenant-specific Backint parameter file for log backups and backups of the backup catalog.

The parameter file must be owned by the operating system user `<sid>adm` and group of the tenant database.

4. Assign the access permissions 640 to the parameter file.

640 allows user `<sid>adm` read and write access to the file; the group has read permission; others have no access permission.

5. Assign the tenant-specific Backint parameter file(s) in the SAP HANA system.

- a. Locate the parameter `data_backup_parameter_file` in the `backup` section of `global.ini`.

For more information, see *Configuring System Properties in SAP HANA Cockpit* or *Configuring System Properties in SAP HANA Studio*.

- b. Open the change dialog.
- c. Specify the new path to the parameter file.
- d. Save.

When you save, the change takes effect immediately.

To change the Backint parameter file for log backups, repeat this procedure for the parameter `log_backup_parameter_file`.

To change the Backint parameter file for backups of the backup catalog, repeat this procedure for the parameter `catalog_backup_parameter_file`.

Alternatively, to change the Backint parameter file setting, you can execute the following SQL statement:

Sample Code

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'DATABASE',
'<database_name>') SET ('backup', '<backup_parameter_file>') =
'<absolute_path_and_name>' WITH RECONFIGURE
```

This statement changes one parameter. If you need to assign a tenant-specific Backint parameter file for both data backups and log backups, you need to execute the statement once for each Backint parameter file.

Example

Assume that you want to assign new Backint parameter files for a tenant database called `<TENANT1>` in an SAP HANA database called `<PR2>`.

With the following statement, you can assign a new parameter file for data backups:

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'DATABASE', 'TENANT1') SET ('backup', 'data_backup_parameter_file') = '/usr/sap/PR2/SYS/global/hdb/opt/config/DB_TENANT1/PR2_TENANT1_data.utl' WITH RECONFIGURE
```

With the following statement, you can assign a new parameter file for log backups:

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'DATABASE', 'TENANT1') SET ('backup', 'log_backup_parameter_file') = '/usr/sap/PR2/SYS/global/hdb/opt/config/DB_TENANT1/PR2_TENANT1_log.utl' WITH RECONFIGURE
```

Related Information

[Configuring System Properties in SAP HANA Cockpit \[page 297\]](#)

[Configuring System Properties in SAP HANA Studio \[page 301\]](#)

[File and Directory Permissions with High Isolation \[page 206\]](#)

10.2.4.3.2 Delta Backups and Third-Party Backup Tools

SAP HANA supports seamless integration of SAP-certified third-party backup tools.

Normally, delta backups will work using the default configuration settings. In some situations, additional steps may be required to create delta backups with a third-party backup tool.

If you are using a third-party tool to create delta backups, consider the following points:

Configuring the Backint Agent

Ensure that the Backint agent executable (**hdbbackint**) is configured correctly.

For more information about configuring your third-party backup tool, see *Configure a Third-Party Backup Tool* and consult the tool vendor documentation.

–l LOG Option

For delta backups, SAP HANA uses the **hdbbackint** level log (option **–l LOG**) in combination with the **hdbbackint** parameter file for data backups. This **hdbbackint** call is sent internally by SAP HANA and is recorded in the `backint.log` file.

⚠ Caution

If a third-party tool uses the option `-l LOG` to specify the log backup container, the log backup container could unintentionally be used for delta backups as well as for log backups. This can potentially cause an error situation.

For this reason, we recommend that you set up two dedicated `hdbbackint` parameter files: one for data backups and one for log backups

If you are in doubt, check with your tool vendor for more details **before** you use delta backups as part of your backup strategy.

Related Information

[Configure a Third-Party Backup Tool \[page 1304\]](#)

10.2.4.3.3 Upgrading a Third-Party Backup Tool

When you upgrade third-party backup software, the following procedure is strongly recommended.

1. Before you start a software upgrade in your production system:
 - Test the upgrade in a test system.
 - Disable automatic log backups through Backint:
Open the *Backup Console* in SAP HANA studio.
Choose the *Configuration* tab and remove the check mark for *Enable Automatic Log Backup*.
Disabling automatic log backup ensures that the backup history is not disrupted.
For more information, see *SAP Note 2009486 (Disable SAP HANA log backups during upgrade of third-party backup tool that supports the Backint for SAP HANA interface)*.
2. While you are upgrading the third-party backup software:
 - Do **not** perform a data backup of the SAP HANA system.
 - Do not make any changes to the backup catalog.
3. After you have upgraded the third-party backup software, enable automatic log backups again.

→ Tip

Monitor the log area to ensure that enough space is available.

If your third-party backup tool is unavailable for an extended period, consider writing the log backups to the file system.

Related Information

[SAP Note 2009486](#)

10.2.4.4 Creating Backups

The sections that follow describe how to create backups.

You can create backups using the following tools:

- SAP HANA cockpit
- SAP HANA studio
- Native SQL
- DBA Cockpit for SAP HANA
DBA Cockpit for SAP HANA is a tool to monitor and administer SAP HANA databases in an ABAP environment.

⚠ Caution

Do not create a full backup after a database fault or other failure has occurred.

Scheduling Backups

You can schedule backups using SAP HANA cockpit or DBA Cockpit for SAP HANA.

To work with DBA Cockpit for SAP HANA, you need a compatible ABAP system.

You can also schedule regular backups using an external scheduler, such as cron.

Related Information

[DBA Cockpit for SAP HANA](#)
[Schedule Backups \[page 1325\]](#)

10.2.4.4.1 Estimate the Space Needed in the File System for a Data Backup

If there is not enough space in the file system for backups, a backup will fail. To ensure that sufficient free space is available, before you back up the database, you should estimate the amount of space that will be needed.

Context

When you back up an SAP HANA database, the estimated backup size is displayed in the backup dialog in SAP HANA cockpit and SAP HANA studio.

This information is read from table `M_BACKUP_SIZE_ESTIMATIONS`.

For more information, see `M_BACKUP_SIZE_ESTIMATIONS` System View in *SAP HANA SQL and System Views Reference*.

To estimate the space required for a backup, make a note of the space requirement from table `M_BACKUP_SIZE_ESTIMATIONS`, and use the SQL statement `BACKUP CHECK` to check that this amount of space is available in the backup destination.

For more information, see *BACKUP CHECK Statement (Backup and Recovery)* in *SAP HANA SQL and System Views Reference*.

i Note

The actual size of a data backup can be larger or smaller than the estimated size.

For example, if data is changed in the database after the size has been estimated and before the backup is created, the actual backup size may be different from the estimated size.

If existing backups are overwritten by backups with the same names, at least **twice the space** in the backup location is needed, because the old backup and the new backup exist for a time in parallel.

It is therefore recommended to keep some additional free space in reserve.

Related Information

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 1317\]](#)

[Create Data Backups and Delta Backups \[page 1314\]](#)

[Temporary Names for File-Based Backups \[page 1271\]](#)

10.2.4.4.2 Create Data Backups and Delta Backups

Using SAP HANA cockpit, you can create data backups and delta backups (differential backups and incremental backups).

Procedure

1. From SAP HANA cockpit, select a SAP HANA database and open the system overview.
2. Choose *Manage database backups*.

The following information from the backup catalog is displayed:

- Time range that the backup catalog covers
- Total size of the backup catalog
- Information about the most recent backups within the time range
This information includes: status, start time, backup type, duration, size, destination type, and a comment

To display more details about a backup, click its row.

3. To customize the information displayed, choose *Filter*.

A dialog box is displayed.

In the dialog box, select or deselect the columns of interest to you. You can filter the following information:

Option	Description
<i>Backup Type</i>	<p>By default, <i>Complete Data Backup</i> and <i>Data Snapshot</i> are selected.</p> <p>Delta backups are included in the backup catalog, but are not displayed by default in SAP HANA cockpit and SAP HANA studio.</p> <div style="border: 1px solid #ccc; padding: 5px;"><p>i Note</p><p>If a delta backup fails, the default filter setting is automatically changed to display delta backups.</p></div> <p>To display information about delta backups, select <i>Differential Backup</i> and <i>Incremental Backup</i>.</p>
<i>Status</i>	<p>You can display backups with the status:</p> <ul style="list-style-type: none"><input type="radio"/> Canceled<input type="radio"/> Failed<input type="radio"/> Prepared <p>A data snapshot has been prepared, but has not been confirmed or abandoned.</p> <div style="border: 1px solid #ccc; padding: 5px;"><p>i Note</p><p>To create a data snapshot, you need to use native SQL.</p></div> <ul style="list-style-type: none"><input type="radio"/> Running<input type="radio"/> Successful
<i>Start Time</i>	You can display backups from a specific time range.

To change the order in which the columns are displayed, choose *Settings*, and use the arrow buttons.

In the same way, you can also customize the backup details pages for each database.

4. Choose *Create Backup*.
5. Specify the backup type:

Option	Description
Complete	A data backup includes all the data structures that are required to recover the database.
Incremental	An incremental backup stores the data changed since the last data backup - either the last data backup or the last delta backup (incremental or differential).
Differential	Differential backups store all the data changed since the last full data backup.

The estimated size of the backup is displayed. This information is read from table `M_BACKUP_SIZE_ESTIMATIONS`.

For more information, see *M_BACKUP_SIZE_ESTIMATIONS System View*.

6. Specify the backup destination type.

Option	Description
File	Writes the backup data to the file system.
Backint	Writes the backup data through a third-party backup tool.

i Note
This option is only available if a third-party backup tool is installed.

The Backint parameters have no effect on the behavior of SAP HANA. For information about the Backint parameters, contact your tool vendor.

- Specify the backup prefix.

The current data and time is proposed by default.

→ Tip

To be able to more easily identify archived backups, it is strongly recommended to use a unique prefix for each backup.

It is recommended to use a timestamp as a unique prefix.

- Specify the backup destination.

Option	Description
For file-based backups:	<p>Ensure that there is sufficient space at the specified backup destination.</p> <p>The default backup destination can be changed as required.</p> <p>More information: <i>Estimate the Space Needed for a Data Backup</i>.</p>
For third-party backup tools:	<p>For third-party backup tools, the destination is always <code>/usr/sap/<SID>/SYS/global/hdb/backint</code>.</p> <p>You can only change the backup prefix.</p>

- To start the backup, choose *Back Up*.

When the backup has been started, the progress of the backup is displayed.

Canceling a Running Backup

You can cancel a running data backup or a delta backup (differential or incremental).

i Note

The option to cancel a backup is only available while a backup is running.

From the backup dialog, choose *Cancel Backup*.

The backup is canceled and you are notified.

After you have canceled a backup, you can start a new data backup.

→ Tip

If you cancel a running backup performed by a third-party backup tool, it is recommended to ensure that any incomplete backups are physically deleted.

i Note

In some situations, it may **not** be possible to cancel a running backup. For example, if it is not possible to access internal locks, or if a running backup. For example, if it is not possible to access internal locks, or if a file cannot be written to an NFS mount.

When all volumes have been backed up, the backup catalog overview is displayed again. Here, you can confirm that the backup was completed successfully.

Related Information

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 1317\]](#)

[Estimate the Space Needed in the File System for a Data Backup \[page 1313\]](#)

10.2.4.4.3 Create Data Backups and Delta Backups (SAP HANA Studio)

Using SAP HANA studio, you can create full data backups (complete data backups and data snapshots) and delta backups (differential backups and incremental backups) of SAP HANA databases.

Prerequisites

- To back up a SAP HANA database, you need the authorizations BACKUP ADMIN and CATALOG READ. To back up a tenant database, you need the authorization DATABASE ADMIN. For more information, see *Authorizations for Backup and Recovery*.
- The database is online, and all configured services are running. To check this in SAP HANA studio: Go to the *Overview* tab and check the *Operational Status*.
- For a file-based backup, there is sufficient space at the backup destination. For more information, see *Estimate the Space Needed in the File System for a Data Backup*.
- For a backup using a third-party tool, the tool is properly configured and connected to the SAP HANA system. For more information, see *Working with Third-Party Backup Tools*.

i Note

With a data backup, the database configuration files (*.ini files) are not backed up. Configuration files that contain customer-specific changes can be backed up manually in order to more easily identify and restore customer-specific changes in a recovery situation.

Procedure

1. Over the *Backup Console*, right-click to open the context menu.

The estimated backup size is displayed.

2. Choose *Back Up System Database...* or *Back Up Tenant Database...*

i Note

A tenant database is backed up through its system database.

The backup dialog appears.

i Note

If you are backing up a tenant database, select the tenant database to be backed up and then choose *Next*.

3. Select a backup type.

Option	Description
Complete Data Backup	A data backup includes all the data structures that are required to recover the database.
Differential Data Backup	Differential backups store all the data changed since the last full data backup.
Incremental Data Backup	An incremental backup stores the data changed since the last full data backup or the last delta backup (incremental or differential).

4. Select a destination type:

Option	Description
File	Writes the backup data to the file system. Each SAP HANA service writes backup data to a separate file in the specified destination in the file system.
Backint	Writes the backup data through a third-party backup tool. Each SAP HANA service starts the <code>Backint for SAP HANA</code> agent and sends the backup data to the third-party backup tool.

i Note

This option is only available if a third-party backup tool is installed.

5. Specify the backup destination.

The default backup destination is the path specified on the *Configuration* tab of the Backup Console.

For file-based backups:

Ensure that there is sufficient space at the specified backup destination.

You can change the default backup destination.

For more information, see *Estimate the Space Needed in the File System for a Data Backup*.

For third-party backup tools:

For third-party backup tools, the destination is always `/usr/sap/<SID>/SYS/global/hdb/backint`.

You can only change the backup prefix.

6. Specify the backup prefix.

→ Tip

To be able to more easily identify archived backups, it is strongly recommended to use a unique prefix for each backup.

It is recommended to use a timestamp as a unique prefix.

For file-based backups, using a unique prefix also prevents existing full data backups from being overwritten in the file system. Delta backups are never overwritten by newer delta backups because SAP HANA asserts a unique file name.

The `Backint for SAP HANA` interface can distinguish between multiple backups with the same name. For this reason, with third-party backup tools, you do not need to use a different prefix for each backup. Nevertheless, for easier identification and versioning, it is strongly recommended to assign unique prefixes to backups created with third-party tools.

7. Choose *Next*.

A summary of the backup settings is displayed.

8. If all settings are correct, choose *Finish*.

The backup is started.

Canceling a Running Backup

You can cancel a running data backup or a delta backup (differential or incremental).

i Note

The option to cancel a backup is only available while a backup is running.

From the backup dialog, choose *Cancel Backup*.

The backup is canceled and you are notified.

After you have canceled a backup, you can start a new data backup.

→ Tip

If you cancel a running backup performed by a third-party backup tool, it is recommended to ensure that any incomplete backups are physically deleted.

i Note

In some situations, it may **not** be possible to cancel a running backup. For example, if it is not possible to access internal locks, or if a running backup. For example, if it is not possible to access internal locks, or if a file cannot be written to an NFS mount.

Results

The backup wizard shows the progress of the backup for all the services.

If you close the backup wizard, you can continue to monitor the progress of the backup on the [Overview](#) tab of the Backup Console.

When all volumes have been backed up, a confirmation message is displayed.

Information about the completed backup is displayed in the Backup Console. Go to ► [Overview](#) ► [Last Successful Data Backup](#) ►.

Related Information

[Authorizations for Backup and Recovery \[page 1244\]](#)

[Estimate the Space Needed in the File System for a Data Backup \[page 1313\]](#)

[Working with Third-Party Backup Tools \[page 1303\]](#)

10.2.4.4.4 Create a Data Snapshot (Native SQL)

You can create a data snapshot using SQL.

Prerequisites

- You can create a data snapshot of:
 - An SAP HANA multitenant database container with **one tenant database**
Currently, a data snapshot of an SAP HANA database with more than one tenant is not supported.
To back up SAP HANA systems with more than one tenant database, use data backups.

i Note

A data snapshot can only be created through the system database.

It is not possible to create a data snapshot for the tenant database separately.

- An SAP HANA single-container system
- The SAP HANA database is online, and all the configured services are running.
To check whether the database is online:

SAP HANA cockpit

On the *Tenant Monitoring and Administration* block, the system status for the system database and the tenant database is `System Running`.

SAP HANA studio

Go to the Administration Console.

Go to the *Overview* tab.

In the section *Operational Status*, all the SAP HANA services should be `started`.

Context

A data snapshot is created in three steps that are performed in the SAP HANA database and at storage system level.

Step to Create a Data Snapshot	Description
<i>Prepare</i> the database for the data snapshot.	An internal database snapshot is created that reflects a consistent database state at the point in time it is created in the file system. i Note If an internal database snapshot exists, no new data backups or new data snapshots can be created. Conversely, while a data backup is running, you cannot create a data snapshot.

Step to Create a Data Snapshot

Description

Create the data snapshot.

The data snapshot is created based on the previously created internal database snapshot.

In the storage system, you need to manually make all the files and directories from the data area available in a separate storage location.

→ Remember

Data snapshots only offer increased data safety if they are moved or replicated to a separate storage medium. The files and directories under the mountpoint of the data area must all be stored together. The data volumes themselves must not be moved.

Confirm or *Abandon* the data snapshot.

If the data snapshot was successfully made available in a new storage location, you can confirm the data snapshot.

A confirm may not always work and could return an error. If the confirm fails, the database snapshot is marked as unsuccessful. You should physically delete the data snapshot because it may not be possible to use it for recovery.

To ensure its consistent state, the **data snapshot** relies on the previously created **internal database snapshot**. If the database, or a database service, is restarted, the **internal database snapshot** is lost.

i Note

The data snapshot is always written - even if the internal database snapshot is lost before the data snapshot is confirmed. During confirmation, you are notified whether the data snapshot can be used for a recovery.

Procedure

To execute SQL statements, you can use the SQL console in SAP HANA cockpit or SAP HANA studio.

1. Create a new internal database snapshot.

Use the following SQL statement:

```
BACKUP DATA FOR FULL SYSTEM CREATE SNAPSHOT [COMMENT <STRING>;
```

❖ Example

```
BACKUP DATA FOR FULL SYSTEM CREATE SNAPSHOT COMMENT 'SNAPSHOT-2017-03-16';
```

Optionally, add a comment. This comment can help to identify the data snapshot in the backup catalog.

For more information, see *BACKUP DATA CREATE SNAPSHOT Statement (Backup and Recovery)* in the *SAP HANA SQL and System Views Reference*.

An internal snapshot is now created.

2. Find out the backup ID of the **internal database snapshot** in the state `PREPARED`.

i Note

SAP HANA cannot ensure that the backup ID of the data snapshot of the system database and the backup ID of the tenant database are not the same. As data snapshots are administered by the system database, you must use the backup ID of the **system database** to create the data snapshot.

Use the following SQL statement:

```
SELECT * FROM M_BACKUP_CATALOG WHERE ENTRY_TYPE_NAME = 'data snapshot';
```

Example

```
SELECT BACKUP_ID, COMMENT FROM M_BACKUP_CATALOG WHERE ENTRY_TYPE_NAME = 'data snapshot' AND STATE_NAME = 'prepared' AND COMMENT = 'SNAPSHOT-2017-03-16';
```

Make a note of the backup ID.

i Note

Older internal database snapshots may exist in the state `successful` or `unsuccessful`.

The database is now prepared for the data snapshot.

An **internal database snapshot** is created, reflecting a consistent database state at the point in time it is created.

i Note

If an internal database snapshot exists, no new data backups or new data snapshots can be created.

Conversely, while a data backup is running, you cannot create a data snapshot.

At this stage, all the snapshot-relevant data is only stored in the data area. To be able to use the data snapshot for a recovery later on, you now need to create the data snapshot. To create a data snapshot, this data needs to be stored in a separate location.

3. In the storage system, make all the files and directories from the data area available together in a separate storage location.

To create the data snapshot, you can use the tool provided by your storage vendor. For more information, consult the tool documentation.

i Note

A data snapshot contains all the persisted data in the data area. For this reason, the files and directories under the mountpoint of the data area must all be stored together.

⚠ Caution

For a recovery using a data snapshot, only the data area must be restored from the storage tool. You still can use the log area for the recovery.

i Note

The directory name of the data area is defined by configuration parameter `basepath_datavolumes` in the `global.ini` configuration file, in the `persistence` section.

After the data snapshot has been created in a separate storage location, it needs to be confirmed.

4. [Confirm](#) or [Abandon](#) the data snapshot.

Use the following SQL statement:

Option	Description
Confirm	<pre>BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID <BACKUP_ID> SUCCESSFUL <STRING>;</pre> <p>Confirm that the data snapshot has been successfully saved to a new storage location.</p> <p>You can specify an external ID to identify the data snapshot later in the storage system.</p>

❖ Example

```
BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID 1489592445498
SUCCESSFUL 'SNAPSHOT-2017-03-16';
```

Abandon	<pre>BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID <BACKUP_ID> UNSUCCESSFUL [<STRING >];</pre>
----------------	---

If the data snapshot cannot be created, or if confirmation fails, choose [Abandon](#).

Optionally, you can add a comment to explain why the data snapshot was not successful.

❖ Example

```
BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID 1489592445498
UNSUCCESSFUL 'SNAPSHOT-2017-03-16 FAILED';
```

For more information, see *BACKUP DATA CLOSE SNAPSHOT Statement (Backup and Recovery)* in the *SAP HANA SQL and System Views Reference*.

→ Tip

It is strongly recommended to confirm or abandon a data snapshot **as soon as possible after it has been created**.

While the data snapshot is being prepared or created, the snapshot-relevant data is frozen. While the snapshot-relevant data remains frozen, changes can still be made in the database. Such changes will not cause the frozen snapshot-relevant data to be changed. Instead, the changes are written to positions in the data area that are separate from the data snapshot. Changes are also written to the log.

However, the longer the snapshot-relevant data is kept frozen, the more the data volume can grow.

i Note

If the database or an individual database service is restarted, the **internal database snapshot** is lost. If the database snapshot is lost before the data snapshot is confirmed, the data snapshot is still written. During confirmation, the database notifies you that the data snapshot cannot be used.

After you have confirmed or abandoned a data snapshot, it is recorded in the backup catalog as either successful or unsuccessful.

i Note

A data snapshot now exists for both the system database and the tenant database.

The internal database snapshot that was used to create the data snapshot is discarded.

It is now possible to create further data snapshots or data backups.

Related Information

[Data Snapshots \[page 1249\]](#)

10.2.4.4.5 Schedule Backups

Using SAP HANA cockpit, you can schedule data backups or delta backups to run at specific intervals.

Prerequisites

- The system privilege BACKUP ADMIN and read authorization for the tables:
 - `_SYS_XS.JOB_SCHEDULES`
 - `_SYS_XS.JOBS`
- Backup schedules must be activated globally.

i Note

The user that activates the backup schedules is also used to execute the backups.

To activate backup schedules from SAP HANA cockpit:

1. From the database overview, choose **Manage database backups** > **Configure Backup**.
2. Go to **Data Backup Settings** > **Data Backup Scheduler**.
3. Choose **Edit**.

4. Set Enable Data Backup Scheduler to YES.
5. [Save](#).

i Note

If it is not possible to enable the data backup scheduler from SAP HANA cockpit, check that the XS Job Scheduler and the XS Engine (XS Classic) are enabled. You may need to manually enable the XS Job Scheduler, and possibly also the XS Engine.

For information about how to manually enable the XS Job Scheduler, see *Enable the XS Job Scheduler (XS Classic)*.

For information about how to manually enable the XS Engine, see *Enable the XS Engine (XS Classic)*.

⚠ Caution

Backup schedules created with SAP HANA cockpit 1.0 are not compatible with SAP HANA cockpit 2.0.

Before you upgrade from SAP HANA 1.0 to SAP HANA 2.0, you must use SAP HANA cockpit 1.0 to delete all the backup schedules created with SAP HANA 1.0.

After the upgrade to SAP HANA 2.0, you need to create new backup schedules.

SAP HANA cockpit 2.0 cannot schedule backups for SAP HANA 1.0 databases.

Procedure

Schedule a Backup

1. From SAP HANA cockpit, open the system overview of the SAP HANA database, for which you want to schedule a backup, and choose [Manage database backups](#).

An overview of the information from the backup catalog is displayed.

2. To display an overview of backup schedules, choose [Go to Schedules](#).

Task	Steps
Display the details for a backup schedule	Click the schedule's row.
Pause or reactivate all backup schedules	Choose Backup Schedules: Off / On .
Pause or activate a specific backup schedule.	Choose Pause to deactivate a backup schedule. Choose Activate to reactivate a paused schedule.
Delete a schedule.	Choose Delete The backup schedule is deleted permanently.

3. To create a new backup schedule, choose [Create Schedule](#).
4. Specify the backup type.

Option	Description
Complete	A complete data backup includes all the data that is required to recover the database to a consistent state.
Incremental	Stores the data changed since the last full data backup (complete data backup or data snapshot) or the last incremental or differential backup.
Differential	Stores all the data changed since the last full data backup (complete data backup or data snapshot).

i Note

Currently, scheduling data snapshots is not supported.

- Specify the *Destination Type*.

Option	Description
<i>File</i>	<p>If you are creating a backup to the file system, select <i>File</i>. If necessary, you can specify a new destination or change the default destination.</p> <p>By default, file-based data backups are written to <code>\$DIR_INSTANCE/backup/data</code>.</p> <p>For more information, see <i>Parameters for Data Backup Settings</i>.</p>
<i>Backint</i>	<p>If you are working with a third-party backup tool, select the destination type <i>Backint</i> and, if needed, specify the <i>Backint parameters</i>.</p> <p>For more information, see <i>Working with Third-Party Backup Tools</i>.</p>

- Specify a *Backup Prefix*.

→ Tip

To be able to more easily identify archived backups, it is strongly recommended to use a unique prefix for each backup.

By default, the name of each scheduled backup is prefixed with the timestamp of the start of the backup. The placeholders [date] and [time] are automatically converted to the current timestamp.

All times specified are interpreted as UTC.

- Optionally, add a comment.

This comment helps you to identify the backups in the backup catalog.

- Specify the settings for the backup schedule.

The *Next Backup* field is completed automatically after you have specified the recurrence intervals.

Schedule Settings	Description
<i>Schedule Name</i>	The name of the schedule.
<i>Start of Schedule</i>	Specify the start time of the schedule.

Schedule Settings	Description
	The first backup in the series will be created after this time.
<i>Recurrence</i>	<ul style="list-style-type: none"> Specify on which days of the week you want a backup to be created. You can select one or more days. Specify the time to create a backup on the selected day(s).

9. Save.

An overview of the backup schedules is displayed.

The new schedule is displayed as *Active*.

The next backup scheduled will run at the first possible time after the start time of the schedule.

i Note

It is not possible to change an existing schedule. If a schedule needs to be changed, you need to delete it and create a new schedule.

⚠ Caution

If SAP HANA is **offline** at a time for which backups are scheduled, scheduled backups will not run.

Note that when SAP HANA is running again, skipped backups are **not automatically rescheduled**.

Related Information

[Enable the XS Job Scheduler \(XS Classic\) \[page 1329\]](#)

[Enable the XS Engine \(XS Classic\) \[page 1330\]](#)

[SAP HANA Backup Types \[page 1246\]](#)

[Parameters for Data Backup Settings \[page 1283\]](#)

[Working with Third-Party Backup Tools \[page 1303\]](#)

10.2.4.4.5.1 Enable the XS Job Scheduler (XS Classic)

To schedule data backups or delta backups SAP HANA cockpit, the XS Job Scheduler (XS Classic) must be enabled.

Prerequisites

The XS Job Scheduler requires the XS Engine to be enabled.

Normally, the XS Engine is enabled by default. In some cases, particularly with SAP HANA Express, the XS Engine is not enabled by default, and you may need to activate the XS Engine manually.

Context

You can normally enable the XS Job Scheduler from SAP HANA cockpit.

For more information, see *Schedule Backups*.

If you need to enable the XS Job Scheduler manually, follow the steps described below.

The XS Job Scheduler is enabled separately for the system database and for each tenant database.

i Note

A backup of a tenant database must be scheduled through the tenant database itself. A backup of a tenant database cannot be scheduled through the system database.

Procedure

1. For the **system database**, the XS Job Scheduler must be enabled in the `nameserver.ini` file.

To enable the XS Job Scheduler, you can use the following SQL statement:

```
ALTER SYSTEM ALTER configuration ('nameserver.ini','SYSTEM') SET ('scheduler','enabled')= 'true' WITH reconfigure;
```

2. For each **tenant database**, the XS Job Scheduler must be enabled in the `xsengine.ini` file.

To enable the XS Job Scheduler, you can use the following SQL statement:

```
ALTER SYSTEM ALTER configuration ('xsengine.ini','SYSTEM') SET ('scheduler','enabled')= 'true' WITH reconfigure;
```

Related Information

[Schedule Backups \[page 1325\]](#)

[Enable the XS Engine \(XS Classic\) \[page 1330\]](#)

10.2.4.4.5.2 Enable the XS Engine (XS Classic)

To schedule backups using SAP HANA cockpit, the XS scheduler must be enabled. The XS scheduler requires the XS engine to be running.

Context

Normally, the XS engine is enabled by default in an SAP HANA system. In some cases, particularly with SAP HANA Express, the XS engine is not enabled by default, and you may need to activate the XS engine manually.

The XS engine is enabled separately for the system database and for each tenant database.

Procedure

1. For the **system database**, the XS engine must be enabled in the `nameserver.ini` file.

To enable the XS engine, you can use the following SQL statements:

```
ALTER SYSTEM ALTER configuration ('nameserver.ini','SYSTEM') SET ('httpserver','embedded')= 'true' WITH reconfigure;
```

```
ALTER SYSTEM ALTER configuration ('nameserver.ini','SYSTEM') SET ('httpserver','workerpoolsize')= '5' WITH reconfigure;
```

→ Tip

It is recommended that you set the workerpool size to 5 for the system database. However, if you need to schedule many backup jobs, consider increasing the value in accordance with your system requirements.

2. For each **tenant database**, the XS engine must be enabled in the `xsengine.ini` file.

To enable the XS engine, you can use the following SQL statement:

```
ALTER SYSTEM ALTER configuration ('xsengine.ini','SYSTEM') SET ('httpserver','embedded')= 'true' WITH reconfigure;
```

```
ALTER SYSTEM ALTER configuration ('xsengine.ini','SYSTEM') SET ('httpserver','workerpoolsize')= '5' WITH reconfigure;
```

→ Tip

It is recommended that you set the workerpool size to 5 for a tenant database. However, if you need to schedule many backup jobs, consider increasing the value in accordance with your system requirements.

Related Information

[Schedule Backups \[page 1325\]](#)

[Enable the XS Job Scheduler \(XS Classic\) \[page 1329\]](#)

10.2.4.5 Backup Audit Actions for Security

You can audit the creation and cancelation of a backup.

When an action occurs, the audit policy is triggered and an audit event is written to the audit trail. Audit policies are database-specific.

Related Information

[Create an Audit Policy \[page 839\]](#)

10.2.5 SAP HANA Recovery

It may be necessary to recover an SAP HANA database due to a number of different reasons.

- Disaster recovery
 - The data area is unusable.
For more information, see *Data Area is Unusable (Disaster Recovery)*.
 - The log area is unusable.
For more information, see *Log Area is Unusable (Disaster Recovery)*.
- Fault recovery
If a logical error occurs, the database needs to be recovered to its state at a particular point in time.
For more information, see *Logical Error – Point-in-Time Recovery (Fault Recovery)*.
- You want to create a copy of the database.
For more information, see *Copying a Database Using Backup and Recovery*.

Related Information

SAP HANA Recovery

[Prerequisites for Database Recovery \[page 1332\]](#)

[Recover a Database \[page 1347\]](#)

[Recover a Database \(SAP HANA Studio\) \[page 1351\]](#)

[Recovering a System Database Using Native SQL \[page 1358\]](#)

[Copying a Database Using Backup and Recovery \[page 1374\]](#)

SAP HANA Recovery Scenarios

[Recovery Scenarios \[page 1368\]](#)

[Data Area is Unusable \(Disaster Recovery\) \[page 1368\]](#)

[Log Area is Unusable \(Disaster Recovery\) \[page 1369\]](#)

[Logical Error – Point-in-Time Recovery \(Fault Recovery\) \[page 1370\]](#)

10.2.5.1 Prerequisites for Database Recovery

Before you start a database recovery, you should ensure that several conditions are fulfilled.

- The SAP HANA database software is installed on the target system.
- You need the following logon credentials:

To Recover...	Log On As...
Tenant database	The database user
System database	The operating system user <sid>adm in the target system

i Note

The SAP control credentials are required for the SAP HANA recovery and also to shut down the database.

To shut down and recover a system, you must first have entered your SAP control credentials for <sid>adm in the resource directory application.

- To perform a recovery, an SAP HANA database needs to be shut down. For this reason, during recovery, a database cannot be accessed by end users or applications.

i Note

If you recover SAP HANA from a data snapshot, you must shut down the database **before** you make the data snapshot available in the data area of the storage system.

For more information, see *Prerequisites: Recovery From a Data Snapshot*.

- The following must be available:
 - At least one full backup (complete data backup or data snapshot) exists.

- If required, delta backups created since the full backup to be used
- If required, log backups created since the full backup to be used (Covering changes not already contained in the delta backups)
- If required, the log area

At the beginning of a recovery, the data backup, the delta backups, and the log backups to be used must be either accessible in the file system or available through a third-party backup tool.

→ Tip

At the beginning of a recovery, SAP HANA checks whether the required data is available.

If you are working with file-based backups, and shared backup storage is not being used, it is not possible to perform these availability checks. For this reason, if recovery-relevant data is not available at the beginning of the recovery, this may not be detected until after the recovery has started. In this situation, the recovery can be started, but will fail.

For this reason, we recommend that you manually check whether a recovery is possible before you start.

For more information, see *Manually Checking Whether a Recovery is Possible*.

i Note

If a full backup is physically available, but not recorded in the backup catalog, that backup can still be used to recover the database. However, it is not possible to recover SAP HANA to a point in time if the log backups or delta backups are not recorded in the backup catalog.

- If you are recovering the database from a data snapshot, the data snapshot must be replicated to the data area.
- To recover **customer-specific configuration settings** (*.ini files), it is recommended that you first configure the customer-specific settings before you recover the database and the replay log backups. For more information, see *Backing Up Customer-Specific Configuration Files*.

Disk Sizing

When you recover an SAP HANA database, disk space requirements can temporarily increase. As a general rule, ensure that you are working with enough disk space to contain at least either the complete data backup and the delta backups or the data snapshot that you are using for the recovery.

Related Information

[Prerequisites: Recovery From a Data Snapshot \[page 1334\]](#)

[Prerequisites: Recovering an Encrypted SAP HANA Database \[page 1334\]](#)

[Prerequisites: Recovery Using Multistreamed Backups \[page 1335\]](#)

[Manually Checking Whether a Recovery is Possible \[page 1335\]](#)

[Points to Note: SAP HANA Recovery \[page 1236\]](#)

10.2.5.1.1 Prerequisites: Recovery From a Data Snapshot

Before you start a database recovery from a data snapshot, you should be aware of several important points.

- To recover SAP HANA from a data snapshot, the data snapshot needs to be made available in the data area of the storage system.

Note

If you recover SAP HANA from a data snapshot, you must shut down the database **before** you make the data snapshot available in the data area of the storage system.

- For a recovery based on a data snapshot, you can optionally also use delta backups and log backups in the same way as with a recovery based on a data backup.
- To recover SAP HANA from a data snapshot, you first need to recover the system database, then the tenant database.
- Data snapshots are only supported for single-tenant systems.

Caution

Using a data snapshot, it is **not possible** to recover an SAP HANA system with more than one tenant database.

If you attempt to recover an SAP HANA system with more than one tenant database from a data snapshot, this may make the data area unusable for all the tenant databases.

Caution

It is **not possible** to use a data snapshot of an **SAP HANA single-container system** to recover an SAP HANA multitenant database container.

10.2.5.1.2 Prerequisites: Recovering an Encrypted SAP HANA Database

Before you recover an encrypted SAP HANA database, you should be aware of several important points.

- If you are recovering SAP HANA from encrypted backups, the backed up encryption root keys must be imported into the instance SSFS.
For more information, see *Import Backed-up Root Keys* in the *SAP HANA Administration Guide (Encryption)*.
- To ensure that all data recovered during data and log recovery is encrypted, data volume encryption must be enabled **before the recovery is started**.
For more information, see *Enable and Disable Encryption of Data and Log Volumes* in the *SAP HANA Administration Guide (Encryption)*.

Related Information

[SAP HANA Backup Encryption \[page 1252\]](#)

[Import Backed-up Root Keys \[page 876\]](#)

[Enabling Encryption of Data and Log Volumes \[page 864\]](#)

[Enable Encryption of Data and Log Backups \[page 872\]](#)

10.2.5.1.3 Prerequisites: Recovery Using Multistreamed Backups

For a recovery using multistreamed backups, there needs to be the same number of channels that were used for the backup.

During a recovery, SAP HANA is able to recognize how many channels were used for a backup, and automatically uses this number of channels for a recovery. SAP HANA does **not** check the value of parameter `parallel_data_backup_backint_channels`.

The backup catalog shows all the parts of a multistreamed backup. For a recovery, the order of the backup parts is not important. SAP HANA can recover the parts of a multistreamed backup in any order.

Related Information

[Multistreaming Data Backups with Third-Party Backup Tools \[page 1306\]](#)

10.2.5.1.4 Manually Checking Whether a Recovery is Possible

The success of a database recovery can only be ensured if the required backups are available and have not been changed since they were created. For this reason, it is recommended that you check backups periodically, or if you suspect that they have been changed in some way since they were created.

When SAP HANA data backups, delta backups, or log backups are created, the integrity of the data to be backed up is automatically checked while the backups are being written. The data is written to the backup destination only if the integrity check was successful.

When a recovery is started, the block-level integrity of the backups to be used is checked automatically. If an error is detected, the recovery is stopped, and will need to be repeated.

In addition to the automatic backup checks performed by SAP HANA, you can manually check data backups and log backups **without performing a recovery**. You can check:

- Whether all the backups needed for a recovery are available and can be accessed
- Whether backups have been changed since they were first written

You can use the following tools to perform manual backup checks:

Tool	What is Checked?
hdbbackupcheck	<p>Checks whether individual data backups and log backups have been changed since they were created.</p> <p>It is recommended that you use this tool periodically to check the consistency of the metadata of a backup.</p>
	<p>⚠ Caution</p> <p>Even if the metadata of a backup is correct, the backup may still have internal errors. For this reason, we recommend that you use <code>hdbbackupcheck</code> to check for corruption in individual data or log backups.</p> <p>For more information, see SAP Note 1869119 (Checking backups with "hdbbackupcheck")</p>
hdbbackupdiag	<p>Determines which data backups and log backups are required to complete a recovery, and also checks whether these backups are available and can be accessed.</p>
	<p>i Note</p> <p>To maintain good recovery performance, and to allow the check to be completed quickly, <code>hdbbackupdiag</code> checks only the metadata of a backup. It does not check the integrity of the backup content on block level.</p> <p><code>hdbbackupdiag</code> can also be used to rebuild the backup catalog. For more information, see <i>Rebuilding the Backup Catalog</i>.</p>

i Note

Both `hdbbackupdiag` and `hdbbackupcheck` can be used with file system backups and third-party backup tools.

`hdbbackupdiag` and `hdbbackupcheck` cannot be used with data snapshots.

With third-party backup tools, `hdbbackupdiag` and `hdbbackupcheck` must run in a system with the same SID to which the backups were written.

i Note

`hdbbackupcheck` does not support SAP HANA Dynamic Tiering.

Isolation Level High for Backups and Third-Party Tools

To grant the system administrator access to the tenant database backup files and directories, you need to add the `<sid>adm` user to the operating system group of each tenant. For more information, see *File and Directory Permissions with High Isolation* in the *SAP HANA Administration Guide*.

Related Information

[Checking Individual Backups \[page 1337\]](#)

[Checking the Backups Required for a Recovery \[page 1342\]](#)

[SAP Note 1869119](#)

[File and Directory Permissions with High Isolation \[page 206\]](#)

[Isolation Level High for Backups and Third-Party Backup Tools \[page 1308\]](#)

[Rebuilding the Backup Catalog \[page 1263\]](#)

10.2.5.1.4.1 Checking Individual Backups

You can use the `hdbbackupcheck` tool to check the integrity of individual data backups and log backups for file-based SAP HANA databases.

Context

You can use the `hdbbackupcheck` tool either from inside an SAP HANA installation, or from outside an SAP HANA installation to check backups that are not accessed by an SAP HANA node.

i Note

Using `hdbbackupcheck` **outside** an SAP HANA installation is recommended only for file-based backups.

`hdbbackupcheck` does not support checking tenant databases with third-party tools.

Related Information

[SAP Note 1869119](#)

10.2.5.1.4.1.1 Check Individual Backups Inside an SAP HANA Installation

You can use the `hdbbackupcheck` tool to manually check the integrity of individual data backups and log backups.

Procedure

A data backup of an SAP HANA instance consists of multiple parts, each with the same prefix. A part of a backup is a backup file in the system storage or a backup object that has been transferred to an external backup tool. To check a data backup, you need to start `hdbbackupcheck` for each individual part of the data backup.

Note

If you are working with third-party tools, consult the tool documentation to learn more about the backup checks that these tools perform.

1. Identify the parts of the data backup that you want to check.

In SAP HANA studio, open the **Backup Console** and go to the **Backup Catalog** tab. Alternatively, use `hdbbackupdiag`.

2. Make a note of the following information:

- File name (*Location*)
For file-based data backups, the location is the file system path to the data backup. If the data backup is below the current directory, the relative path can be used.
For data backups managed using a third-party backup tool, the location is the complete path and name, beginning with `/usr/sap/<SID>/SYS/global/hdb/backupint/`.
- External backup ID
You need the external ID if you are using a third-party backup tool.
- Optionally, the backup ID assigned by the SAP HANA database when the backup was created.

3. Call `hdbbackupcheck` using the appropriate values for each part of a data backup.

To start `hdbbackupcheck` on the command line, use the following options:

```
hdbbackupcheck [parameters] <backup> [-i <backupid>] [-e <ebid>]
```

Options for `hdbbackupcheck`

Option	Description
<code>-v</code>	Display the header data of the backup.
<code>-p <directory></code>	By default, the log files <code>backupcheck.log</code> and <code>backintcheck.log</code> are created in the trace directory. To create the log files in a different directory, call <code>hdbbackupcheck</code> with option <code>-p <directory></code> and specify the directory.

Option	Description
<code><backup></code>	Name of the backup file.
<code>--backintParamFile <filename></code>	Specify the parameter file for the third-party backup tool. If the working directory is not the directory where the file is located, specify the absolute path. To find out this path, consult the documentation provided by the tool vendor.
<code>-i <backupid></code>	Specify the SAP HANA backup ID of the backup to be checked. The backup ID is assigned to the backup when it is created.
<code>-e <ebid></code>	External backup ID If the part of the backup is in a third-party backup tool, you need to specify the external backup ID.
<code>--dump <backupfile></code>	List the contents of the backup, if possible.

Results

hdbbackupcheck notifies you if any errors were detected in the checked part of the backup.

10.2.5.1.4.1.1.1 Examples of Output From hdbbackupcheck

hdbbackupcheck notifies you if any errors were detected in the checked part of the backup.

If no errors were detected, hdbbackupcheck returns 0.

If an error was detected, hdbbackupcheck returns 1.

Below are some examples of output from hdbbackupcheck:

❁ Example

```
hdbbackupcheck backup/data/BackupTestMaster_databackup_1_1
```

⇐ Output Code

```
Backup '/hana/shared/BNR/HDB00/backup/data/BackupTestMaster_databackup_1_1'  
successfully checked.
```

❁ Example

```
hdbbackupcheck -v backup/data/BackupTestMaster_databackup_1_1
```

⌘ Output Code

```
Check backup '/hana/shared/BNR/HDB00/backup/data/BackupTestMaster_databackup_1_1'
Check backup '/hana/shared/BNR/HDB00/backup/data/BackupTestMaster_databackup_1_1'.
Destination header information:
DestVersion: 5
DatabaselD: 51a3a622-1627-46c8-e100-00000a1d0eab
InternalStartTime: 1370415876795
CurrDestInformation: [FILE][/usr/sap/BNR/HDB00/backup/data/
BackupTestMaster_databackup_1_1]
backupID: 1370415876776
ServiceName: nameserver
NumberOfVolumeFiles: 4
HostName: berl30052174a
VolumelD: 1
DestID: 1
MaxDestID: 1
SrcPoolInformation[0]: [DATABASE_SNAPSHOT]@node[1]
DstPoolInformation[0]: [FILE][/usr/sap/BNR/HDB00/backup/data/
BackupTestMaster_databackup_1_1]
Source header information:
SrcType: 1
SourceInformation: [DATABASE_SNAPSHOT]@node[1]
srcVersion: 5
sourceSize: 70455296
```

❁ Example

```
hdbbackupcheck backup/data/Hallo_databackup_2_1
```

Checking a data backup that was written to the file system and at some stage was corrupted:

⌘ Output Code

```
ERROR: [110088] Error reading backup from 'FILE' '/hana/shared/BNR/HDB00/backup/data/
Hallo_databackup_2_1'
ERROR: [110059] The backup /hana/shared/BNR/HDB00/backup/data/Hallo_databackup_2_1 is
corrupt, size is 14807859
```

ERROR: Backup '/hana/shared/BNR/HDB00/backup/data/Hallo_databackup_2_1' not successfully checked!

❖ Example

Checking a log backup that was saved to a third-party backup tool using the configuration file /myBackupTool/backupint.cfg:

```
hdbbackupcheck --backupintParamFile /myBackupTool/backupint.cfg /usr/sap/BNR/SYS/global/hdb/backupint/log_backup_1_0_2177088_2177344 -e BCKINTk168Gc
```

↵ Output Code

Backup '/usr/sap/TG2/SYS/global/hdb/backupint/log_backup_1_0_2177088_2177344' successfully checked.

10.2.5.1.4.1.2 Check Individual Backups Outside an SAP HANA Installation

You can use the `hdbbackupcheck` tool to manually check the integrity of backups **outside** an SAP HANA installation.

Context

You can use `hdbbackupcheck` to check backups that are not accessed by an SAP HANA node, without generating additional load on the SAP HANA node.

i Note

Using `hdbbackupcheck` **outside** an SAP HANA installation is recommended only for file-based backups.

Procedure

1. In the SAP HANA installation, create an archive with the required files: `hdbbackupcheckpack <archive>`

i Note

The archive created here contains only the test software and not the data backups to be tested.

2. Move the archive `<archive>` to the target system and unpack it:
 - a. Create a directory `<targetdir>` in the target system.

- b. Copy the archive `<archive>` and the program `$DIR_INSTANCE/exe/SAPCAR` to the directory `<targetdir>` in the target system.
3. Unpack the archive, add the directory of `hdbbackupcheck` to the environment variable `LD_LIBRARY_PATH`, and call the program as described above:
 - a. `cd <targetdir>`
 - b. `./SAPCAR -xvf <archive>`
 - c. `export LD_LIBRARY_PATH=<targetdir>:$LD_LIBRARY_PATH`
 - d. `./hdbbackupcheck -v <backup>`

i Note

By default, the files `backupcheck.log` and `backintcheck.log` are created in the current directory. To create these files in a different directory, start `hdbbackupcheck` with option `-p <directory>`.

10.2.5.1.4.2 Checking the Backups Required for a Recovery

The `hdbbackupdiag` tool determines which backups are required to complete a recovery to a specified point in time. It also checks whether these backups are available and whether they can be accessed.

Context

With `hdbbackupdiag`, you can verify the following:

For file-based backups:

- The backup is available in the file system, either at the location to which it was written or at a location specified by a search path.
The backups to be used can be in any directory in the file system.
- The current operating system user has read authorization for the file.
- The actual size of the backup file is the same as the size recorded in the backup file header.
- The backup ID is identical to the backup ID specified in the backup catalog.

For backups created using a third-party backup tool:

- The backup is available in the third-party backup tool.

Procedure

1. Ensure that `hdbbackupdiag` can locate the backup catalog.

To do this, you can either execute `hdbbackupdiag` in the directory where the backup catalog is located (file-based only) or execute `hdbbackupdiag` with options that specify which directories to search for the latest backup catalog. By default, this is the directory where the last log backups were written before the recovery was started.

The default directory is `$DIR_INSTANCE/backup/log`.

For more information, see *Change the Log Backup Settings*.

2. Start `hdbbackupdiag`.

Using the following command:

```
hdbbackupdiag [options] [-d <directory>]
```

The options for `hdbbackupdiag` are described below.

Option	Description
<code>-h --help</code>	Display the available options.
<code>--check</code>	Check whether the metadata is correct and consistent and has not changed since the backup was made.
<code>-f</code>	Display the names of the backups required for recovery as a simple list. From this list, the backup names can be easily included in shell scripts.
<code>-B</code>	Display the names of backups with Backint information.
<code>-v</code>	Display all available information. For example, the SAP HANA version that was used to create a backup.
<code>-d <directory></code>	Specify the directory to search for the backup catalog. If you do not specify a directory, the current directory is searched for the latest version of the backup catalog. If specified, the directories indicated with <code>--logDirs</code> (see below) and the third-party backup tool are also searched.
i Note All directories must be specified as absolute paths.	
<code>-c <catalog></code>	Specify the name of the backup catalog.
<code>-i <BackupID></code>	Specify a backup ID. If you do not specify a backup ID, the most recent usable data backup is used.
<code>-u <"YYYY-MM-DD hh:mm:ss"></code>	Specify a target time for the recovery (UTC time). If you do not specify a time, the most recent possible point in time is used.
<code>--dataDir <directory></code>	Specify a directory to search for data backups or delta backups. If you do not specify a directory, only the paths specified in the backup catalog are searched.
<code>--logDirs <directories></code>	Specify a comma-separated list of directories to search for log backup files. If you do not specify this option, only the paths specified in the backup catalog are searched.

Option	Description
<code>--useBackintForCatalog</code>	With this option, the third-party backup tool is searched for the most recent version of the backup catalog.
<code>--databaseName <database></code>	Used only with SAP HANA and third-party backup tools. This option is used with <code>--useBackintForCatalog</code> to specify a tenant database or the system database: <code>--databaseName <name_of_tenant_database></code> <code>--databaseName SYSTEMDB</code>
<code>--backintDataParamFile <paramFileName></code>	Specify a parameter file to access data backups and delta backups through a third-party backup tool.
<code>--backintLogParamFile <paramFileName></code>	Specify a parameter file to access log backups through a third-party backup tool. If you do not specify a parameter file, the parameter file used to access the data backups is used.
<code>--pickCatalog</code>	If you want to recover SAP HANA to a point in time (UNTIL timestamp) that is not available in the current timeline, a suitable catalog is selected for the recovery time. More information: See SAP Note: 2050606 (Recover database from not current backup history)
<code>--generate</code>	Generate a new backup catalog
	<div style="border: 1px solid #ccc; padding: 10px; background-color: #f9f9f9;"> <p>⚠ Caution</p> <p>The <code>--generate</code> option is intended for exceptional situations and file-based backups only. This option does not support data snapshots or backups created using third-party tools.</p> <p>If you use the <code>--generate</code> option, any information about data backups or log backups created using third-party tools will be lost from the newly generated backup catalog. It will then no longer be possible to use this information for a recovery.</p> </div>
<code>--ignoreDeltaDataBackups</code>	Exclude delta backups.

Results

The output of `hdbbackupdiag` contains the names of all the files required to recover the SAP HANA database.

Related Information

[Change the Log Backup Destination Type \[page 1294\]](#)

10.2.5.1.4.2.1 Examples of Output From hdbbackupdiag

The output of `hdbbackupdiag` contains the names of all the files required to recover a SAP HANA database.

If you specify the `--check` option, the results of the metadata checks are also displayed.

❁ Example

```
lu059113:/usr/sap/HD2/HDB00/backup> hdbbackupdiag --check --logDirs /usr/sap/HD2/HDB00/backup/log/ --dataDir /usr/sap/HD2/HDB00/backup/data/
```

This example does the following:

- The directory `/usr/sap/HD2/HDB00/backup/data/` is searched for data backups and delta backups.
- The directory `/usr/sap/HD2/HDB00/backup/log` is searched for log backups.
- The metadata of the backup files is checked to determine whether all the required backups are available and consistent, and whether a recovery to the desired point in time is possible.
- SAP HANA decides which data backup in the specified directory to use for the recovery.

This command yields the following output, including some errors:

📄 Sample Code

```
found backup catalog 1426152872410 from backint /usr/sap/HD2/SYS/global/hdb/backint/log_backup_0_0_0_0
found backup catalog 1426152780165 from file /usr/sap/HD2/HDB00/backup/log/log_backup_0_0_0_0.1426152780165
using backup catalog 1426152872410 from backint /usr/sap/HD2/SYS/global/hdb/backint/log_backup_0_0_0_0
Backup '/usr/sap/HD2/HDB00/backup/data/COMPLETE_DATA_BACKUP_databackup_0_1' successfully checked.
Backup '/usr/sap/HD2/HDB00/backup/data/COMPLETE_DATA_BACKUP_databackup_1_1' successfully checked.
Backup '/usr/sap/HD2/HDB00/backup/data/COMPLETE_DATA_BACKUP_databackup_2_1' successfully checked.
Backup '/usr/sap/HD2/HDB00/backup/data/COMPLETE_DATA_BACKUP_databackup_3_1' successfully checked.
Backup '/usr/sap/HD2/HDB00/backup/log/log_backup_1_0_380224_384576.1426152365477' successfully checked.
Backup '/usr/sap/HD2/HDB00/backup/log/log_backup_1_0_384576_385216.1426152395479' successfully checked.
Backup '/usr/sap/HD2/HDB00/backup/log/log_backup_1_0_385216_385984.1426152425481' successfully checked.
ERROR: [111119] file '/usr/sap/HD2/HDB00/backup/log/log_backup_1_0_385984_386816.1426152455484' not found
ERROR: Backup '/usr/sap/HD2/HDB00/backup/log/log_backup_1_0_385984_386816.1426152455484' check failed.
(...)
Backup '/usr/sap/HD2/HDB00/backup/log/log_backup_2_0_265408_271040.1426152365475' successfully checked.
Backup '/usr/sap/HD2/HDB00/backup/log/log_backup_2_0_271040_271104.1426152396044' successfully checked.
```

```

Backup '/usr/sap/HD2/HDB00/backup/log/
log_backup_2_0_271104_271680.1426152440145' successfully checked.
ERROR: Backup '/usr/sap/HD2/HDB00/backup/log/
log_backup_2_0_271680_272576.1426152500624' has size 45131 bytes, but is
expected to be at least 61440 bytes
ERROR: [110059] The backup /usr/sap/HD2/HDB00/backup/log/
log_backup_2_0_271680_272576.1426152500624 is corrupt, size is 45131 bytes
ERROR: Backup '/usr/sap/HD2/HDB00/backup/log/
log_backup_2_0_271680_272576.1426152500624' check failed.
Backup '/usr/sap/HD2/HDB00/backup/log/
log_backup_2_0_272576_273152.1426152560136' successfully checked.
Backup '/usr/sap/HD2/HDB00/backup/log/
log_backup_2_0_273152_273792.1426152620143' successfully checked.
(...)

```

The first error occurs because a log backup is not available. The second error is because a log backup does not have the expected size.

❁ Example

Display all the backups required to recover the database until May 11, 2015, 01:05:00 p.m. The metadata of the backups is not checked.

From this list, the backup names can be easily included in shell scripts.

The time specified is UTC time, not local time.

```
hdbbackupdiag -f -d /usr/sap/MBY/HDB01/backup/log -u "2015-05-11 13:05:00"
```

```

2015-05-11_13-05_databackup_0_1
2015-05-11_13-05_databackup_1_1
2015-05-11_13-05_databackup_2_1
2015-05-11_13-05_databackup_3_1
2015-05-11_13-05_databackup_4_1
log_backup_1_0_426304_427776.1431332400000
log_backup_1_0_427776_428288.1431333300000
log_backup_1_0_428288_428608.1431334200000
log_backup_1_0_428608_429376.1431335100000
log_backup_1_0_429376_429696.1431336000000
log_backup_1_0_429696_430464.1431336900000
log_backup_1_0_430464_430784.1431337800000
log_backup_1_0_430784_431552.1431338700000
log_backup_1_0_431552_431872.1431339600000
log_backup_2_0_598080_598592.1431340500000
log_backup_2_0_598592_598976.1431341400000
log_backup_2_0_598976_599360.1431342300000
log_backup_2_0_599360_602304.1431343200000
log_backup_2_0_602304_602688.1431344100000
log_backup_3_0_536064_538304.1431345000000
log_backup_4_0_1190656_1191360.1431345900000
log_backup_4_0_1191360_1191680.1431346800000
log_backup_4_0_1191680_1192000.1431347700000
log_backup_4_0_1192000_1192640.1431348600000
log_backup_4_0_1192640_1192960.1431349500000

```

10.2.5.2 Recovering an SAP HANA Database

SAP HANA supports different options for database recovery.

You can recover an SAP HANA database to its most recent consistent state or to an earlier state. You can recover an SAP HANA database to the same system, or to a different system to create a copy of the database.

i Note

To recover a database, it is possible to use a combination of backups from a third-party backup tool and backups from the file system, provided that the backups originate from the same SAP HANA database.

To copy a database, it is not possible to mix backups from the different sources. The backup catalog, the data backups, and the log backups must be from either **only** a third-party backup tool or **only** the file system.

Related Information

[Recover a Database \[page 1347\]](#)

[Recover a Database \(SAP HANA Studio\) \[page 1351\]](#)

[Recover a Database to a Specific Data Backup \(SAP HANA Studio\) \[page 1355\]](#)

[Recovering a System Database Using Native SQL \[page 1358\]](#)

[Recovery Scenarios \[page 1368\]](#)

[Copying a Database Using Backup and Recovery \[page 1374\]](#)

10.2.5.2.1 Recover a Database

Using SAP HANA cockpit, you can recover an SAP HANA database to its most recent state or to a specific point in time.

Context

A tenant database is recovered through its system database.

A tenant database can be recovered to its most recent state or to a specific point in time.

For a system database, a point-in-time recovery is not possible using SAP HANA cockpit. To recover a system database to a point in time, use SQL.

For more information, see *Recovering a Database Using Native SQL*.

i Note

A recovery to the most recent state or to a point in time is equivalent to the SQL statement `RECOVER DATABASE (not RECOVER DATA)`.

For more information, see *RECOVER DATABASE Statement (Backup and Recovery)* and *RECOVER DATA Statement (Backup and Recovery)* in the *SAP HANA SQL and System Views Reference*.

Procedure

1. From SAP HANA cockpit, select the system database.

Option	Description
To recover the system database	Choose <i>Recover database</i> .
To recover a tenant database	Choose <i>Overall Tenant Statuses</i> . Select the tenant database from the overview and choose <i>Recover tenant</i> . If the database is not already shut down, you are prompted to shut it down.

Follow the steps on the screen.

2. When the database is shut down, specify the state to which you want to recover the database.

Option	Description
<i>Recover to the most recent state</i>	Recovers the database to a state as close as possible to the current time. → Tip Using the most recent available full backup makes for a faster recovery.
<i>Recover to a specific point in time</i>	Specify a time zone and a point in time to which to recover the tenant database. i Note Any changes that were made after the specified point in time will not be in the recovered tenant database. i Note The time specified is in UTC. If you specify a point in time in the future, the effect will be the same as recovering the database to the most recent state.

3. Proceed to the next step.
4. Specify the location of the most recent backup catalog.

Option	Description
<i>Backint location only</i>	If a third-party backup tool is selected, Backint is searched.

Option	Description
<i>Default location</i>	<p>For file system backups, the default location for the backups of the backup catalog for the system database is defined using the parameter <code>basepath_catalogbackup</code>.</p> <p>The default setting for <code>basepath_catalogbackup</code> is:</p> <pre>\$DIR_INSTANCE/backup/log</pre> <p>Log backups for tenant databases are written by default to a tenant-specific subdirectory.</p> <p>Backups of the backup catalog are written by default to the same tenant-specific subdirectory as the log backups.</p>
<i>Alternative location</i>	<p>If the backup catalog is not in the default location, specify its location.</p> <p>For more information, see <i>Destination for Backups of the Backup Catalog</i>.</p>

An overview of available backups is displayed.

5. Proceed to the next step.
6. Select the data backup to use for the recovery.
7. Proceed to the next step.
8. Specify whether to use delta backups (differential or incremental backups).

⚠ Caution

The full backups and the delta backups must be in the same location for the recovery to work correctly.

For more information about delta backups, see *Delta Backups*.

9. Proceed to the next step.
10. If necessary, you can change the location to be searched for data backups and delta backups.

To specify additional locations for log backups, choose *Add more*.

If you leave the locations empty, SAP HANA uses the backup locations specified in the backup catalog.

11. Proceed to the next step.
12. Check whether the backups are available.

Here, you can decide whether to check if all the backups needed are available and can be accessed **before the recovery starts**. The availability check is performed at the beginning of the recovery.

i Note

SAP HANA does not check the integrity of the backups content on block level.

For more information, see *Manually Checking Whether a Recovery is Possible*.

13. Proceed to the next step.
14. Specify whether to initialize the log area.

⚠ Caution

If you initialize the log area, the content of the log area is lost.

No log entries from the log area can then be replayed during the recovery.

You **must** initialize the log area in the following situations:

- The log area is unusable.
- You are recovering the database to a different system.

15. Choose [Review](#).

An overview of the settings for the recovery is displayed.

To change any settings, choose [Edit](#).

All the settings that you specified are retained until you change them.

16. To display the SQL statement to be used for the recovery, choose [Display SQL Statement](#).

For more information, see [Recovering a Database Using Native SQL](#).

17. To perform the recovery, choose [Start Recovery](#).

The progress of the recovery for each SAP HANA service is displayed.

Results

When the recovery is completed, a message confirms this, and shows the point in time to which the database was recovered.

i Note

The SQL statement used for a recovery is recorded in `backup.log`. For a point-in-time recovery, the point in time is specified in the SQL statement as **UTC**.

The time at which the recovery was started and completed is recorded in `backup.log` as **local server time**, not UTC.

For more information, see [Diagnosis Files for Backup and Recovery](#).

i Note

The point in time that SAP HANA returns after a recovery may be **before** the point in time that you specified for the recovery. This is because the point in time that was actually reached in the recovery is that of the most recent global COMMIT to the database that was recovered.

The SAP HANA database is now online and can be used by applications.

Related Information

[Prerequisites for Database Recovery \[page 1332\]](#)

[Create a Data Snapshot \(Native SQL\) \[page 1320\]](#)

[Destination for Backups of the Backup Catalog \[page 1299\]](#)

[Delta Backups \[page 1247\]](#)

[Manually Checking Whether a Recovery is Possible \[page 1335\]](#)

[Cancel a Recovery \[page 1364\]](#)

[Recovering a System Database Using Native SQL \[page 1358\]](#)

[Diagnosis Files for Backup and Recovery \[page 1272\]](#)

10.2.5.2.2 Recover a Database (SAP HANA Studio)

Using SAP HANA studio, you can recover an SAP HANA database to its most recent consistent state or to a specific point in time.

Prerequisites

- A data backup or a data snapshot.
- Delta backups and/or log backups
- The log area

For more information, see *Prerequisites for Database Recovery*.

Procedure

1. In SAP HANA studio, open the context menu and choose *Backup and Recovery*.
Then choose either *Recover System Database...* or *Recover Tenant Database...*
The Recovery Wizard opens.
Follow the instructions in the wizard.
2. If prompted, enter the `<sid>adm` user and password and choose *OK*.

i Note

`<sid>adm` is the OS user for the system database, and is not needed for tenant databases.

3. When prompted, confirm that the database can be shut down.
4. Specify the recovery type.

Option	Description
<i>Recover the database to its most recent state</i>	Recovers the database to as close as possible to the current time.

→ Tip

Using the most recent available full backup makes for a faster recovery.

Option	Description
Recover the database to the following point in time	<p>Specify a point in time to which to recover the database.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p>i Note</p> <p>Any changes that were made after the most recent log backup will not be in the recovered database. In addition, all the transactions that were open during the log backup will be rolled back.</p> <p>If you need to perform a point-in-time recovery, consider recovering the database to a different system.</p> </div> <p>This option uses the following data:</p> <ul style="list-style-type: none"> ○ Recommended: The last available data backup (file-based or Backint) or data snapshot available before the specified point in time ○ Log backups and delta backups made after the full backup and up to the desired point in time ○ Log area <div style="background-color: #f0f0f0; padding: 10px;"> <p>i Note</p> <p>If you specify a point in time in the future, the effect will be the same as recovering the database to the most recent state.</p> </div>
<p>▶ Advanced ▶</p> <p>Recover the database to the following log position ▶</p>	<div style="background-color: #f0f0f0; padding: 10px;"> <p>i Note</p> <p>This recovery type is an advanced option that can be used in exceptional cases where a previous recovery failed.</p> </div> <p>This option uses the following data:</p> <ul style="list-style-type: none"> ○ The most recent data backup or data snapshot available before the specified log position ○ Log backups or delta backups made since the data backup or data snapshot to be used ○ The log area <div style="background-color: #f0f0f0; padding: 10px;"> <p>i Note</p> <p>Currently, this option is not available for system databases.</p> </div>

5. Choose [Next](#).
6. If the log backups are not in the original location, specify a new location, and choose [Add](#).
7. Choose [Next](#).

An overview of data backups is displayed.

8. From the backup catalog, you can select a complete data backup or a data snapshot. You can also resume an interrupted recovery.

Option	Description
Refresh	If you make available a new data snapshot in the data area, and it does not immediately appear, choose Refresh to update the overview.
Show More	To display additional backups from the backup catalog, choose Show More .

Option	Description
Check Availability	<p>To ensure that a backup exists at the specified location, choose Check Availability.</p> <p>If the system indicates that the data backup is not available at the selected location, and you know that it has been moved, you can specify an alternative location to be checked.</p>

i Note

The availability of Backint backups can be checked.

The availability of file-based backups can only be checked if shared backup data is being used.

A partially completed recovery that can be resumed is given the backup prefix [RESUME](#).

Caution

The full backups and the delta backups must be in the same location for the recovery to work correctly.

i Note

To recover a database from a data snapshot, the data snapshot must be made available in the data area of the database.

After the system database has been recovered from a data snapshot, the tenant database must then be recovered.

9. Choose [Next](#).
10. Finalize the recovery settings.

Option	Description
Check Availability of Delta and Log Backups	<p>Check whether all the required log backups and delta backups are available before the recovery starts. If any log backups are missing, they are listed, and the recovery is stopped before any data is changed.</p> <p>You can check the availability of either file-based, third-party backups (Backint), or both.</p>

 Caution

If you choose not to perform this check before the recovery starts, the check is still performed, but later in the recovery process.

If an error is not detected until after the recovery has been started, the recovery will be interrupted.

After a recovery has been interrupted, the database has an inconsistent state, and it will not be possible to start the database. If the database has an inconsistent state, SAP HANA automatically prevents the database from starting.

If you attempt to restart the database after a recovery has been interrupted, the following message is written to the nameserver trace file:

Option	Description
	<p>Cannot start the service 'nameserver' at '<host:SQL Port>' responsible for the volume '<volume number>' because of an error during recovery.</p> <p>In this situation, you need to recover the database using a different recovery strategy.</p> <p>i Note</p> <p>Shared backup storage</p> <p>If you are working with file-based backups, and shared storage is not used for backups, the master name server has no access to the backup storage of the other servers. As a consequence, the master name server cannot check whether backups are available. This means that the availability checks cannot be performed at the beginning of the recovery. If you have started a recovery that cannot be completed because one or more of the required backups is not available, this will only be detected later, when each service checks the availability of its own backups.</p> <p>If the complete recovery needs to be repeated because log backups or delta backups are missing, this may cause significant disruption to work with the database.</p>
Initialize Log Area	<p>If you initialize the log area, the content of the log area is lost. No log entries from the log area can then be replayed. The log entries from the log backups are replayed if they are needed.</p> <p>⚠ Caution</p> <p>Disabling log backups may cause significant loss of data.</p> <p>You must select the <i>Initialize log area</i> option in the following situations:</p> <ul style="list-style-type: none"> ○ The log area is unusable ○ You are recovering the database to a different system
Use Delta Backups	<p>By default, SAP HANA includes delta backups in its recovery strategy, and gives preference to delta backups over log backups.</p> <p>You can choose to not use delta backups for a recovery. If delta backups are not used, log backups will be used.</p>
Install New License Key	<p>If you already have a license key for the new SAP HANA database, you can import your existing license key.</p> <p>If you are recovering the database to a database with a new SID or landscape ID, a new license key is needed.</p> <p>For more information, see <i>Points to Note: License Key and Recovery</i>.</p>

11. Choose *Next*.

Option	Description
Show SQL Statement	Display the SQL statement to be used for the recovery.

A summary of the selected options is displayed. To make changes, choose *Back*.

12. If the settings are correct, choose *Finish*.

The recovery starts.

Results

When the recovery is complete, a message confirms this, and shows the timestamp to which the recovery was completed.

i Note

The timestamp that SAP HANA returns after a recovery is the timestamp of the last COMMIT to the database that has been recovered.

This timestamp may be before the point in time that you specified for the recovery.

The SAP HANA database is now online and can be used by applications.

Related Information

[Prerequisites for Database Recovery \[page 1332\]](#)

[Checking the Backups Required for a Recovery \[page 1342\]](#)

[Points to Note: License Key and Recovery \[page 1238\]](#)

[Recovery Scenarios \[page 1368\]](#)

[Copying a Database Using Backup and Recovery \[page 1374\]](#)

[Recovering a System Database Using Native SQL \[page 1358\]](#)

[Cancel a Recovery \[page 1364\]](#)

10.2.5.2.3 Recover a Database to a Specific Data Backup (SAP HANA Studio)

Using SAP HANA studio, you can recover an SAP HANA database from a specific data backup.

Prerequisites

A data backup or a data snapshot.

For more information, see *Prerequisites for Database Recovery*.

Context

i Note

For a recovery from a data backup, delta backups are not used.

Delta backups are only used to recover SAP HANA to a point in time.

i Note

If you are not using the backup catalog for the recovery, you need to know the backup type (File, Backint, or data snapshot), the location, and the prefix of the data backup.

Procedure

1. In SAP HANA studio, open the context menu for a database.
To recover a tenant database, open the context menu from its system database.
2. Choose *Backup and Recovery*.
3. Choose *Recover Tenant Database...* and specify the tenant database to recover.

The Recovery Wizard opens.

Follow the on-screen instructions.

4. Choose *Next*.
5. Specify the following recovery type: *Recover the database to a specific data backup*

⚠ Caution

With this option, the SAP HANA database is **initialized** with the specified data backup. This data backup begins a new database lifecycle. Older data backups are then no longer compatible with logs written after the recovery.

⚠ Caution

Log entries are not replayed from the log backups nor from the log area.

All the log entries that still exist in the log area are deleted.

All the changes made after the data backup or data snapshot will be lost.

If you recover a system database, all the changes to the information about its tenant databases that were made after the data backup or data snapshot will be lost.

6. Choose *Next*.
7. Specify the location of the backup catalog(s).
 - *Recover using the backup catalog*
If the backup catalog is not in the default location, specify the new location.
 - *Recover without the backup catalog*

If you use a backup that is not recorded in the backup catalog, you will need to manually specify the backup type (File or Backint), the location of the backup, and its prefix.

i Note

The SID of the source system is only relevant for database copy using third-party backup tools.

8. Choose *Next*.
9. If prompted, confirm that the database can be shut down.

An overview of the relevant full backups is displayed.

10. Specify the data backup to use for the recovery.

Option	Description
<i>Refresh</i>	If you make available a new data snapshot in the data area, and it does not immediately appear, choose <i>Refresh</i> to update the overview.
<i>Show More</i>	To display additional backups from the backup catalog, choose <i>Show More</i> .
<i>Check Availability</i>	To ensure that a backup exists at the specified location, choose <i>Check Availability</i> . If the system indicates that the data backup is not available at the selected location, and you know that it has been moved, you can specify an alternative location to be checked.

i Note

The availability of Backint backups can be checked.

The availability of file-based backups can only be checked if shared backup data is being used.

11. Choose *Next*.
12. Finalize the recovery settings.

Option	Description
Initialize Log Area	This option is selected by default.
	<div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>⚠ Caution</p> <p>No log entries from the log area are replayed, and the log area is initialized. The content of the log area is lost.</p> </div>
Install New License Key	If you already have a license key for the new SAP HANA database, you can import your existing license key. If you are recovering the database to a database with a new SID or landscape ID, a new license key is needed. More information: <i>Points to Note: License Key and Recovery in Related Information</i>

13. Choose *Next*.

A summary of the selected options is displayed. To make changes, choose *Back*.

To display the SQL statement to be used for the recovery, choose *Show SQL Statement*.

14. If the settings are correct, choose *Finish* to start the recovery.

The progress of the recovery for each SAP HANA service is displayed in the dialog box.

Results

When the recovery is complete, a message confirms this and shows the timestamp when the recovery was completed.

When the SAP HANA database has been restarted, it can be used again by applications.

Related Information

[Authorizations for Backup and Recovery \[page 1244\]](#)

[Prerequisites for Database Recovery \[page 1332\]](#)

[Checking the Backups Required for a Recovery \[page 1342\]](#)

[Points to Note: License Key and Recovery \[page 1238\]](#)

[Starting and Stopping Distributed SAP HANA Systems Using SAPControl \[page 1448\]](#)

[Cancel a Recovery \[page 1364\]](#)

10.2.5.2.4 Recovering a System Database Using Native SQL

To recover an SAP HANA database, it is strongly recommended that you use SAP HANA cockpit. To recover a system database (or an SAP HANA single-container system) using SQL, you can use the `recoverSys.py` tool.

Prerequisites

- The system database is offline.
- You are logged with the OS user `<sid>adm`.

Procedure

- To call `recoverSys.py`, enter the statement in the following format: `HDBSettings.sh recoverSys.py [<parameters>]`

If you run `HDBSettings.sh recoverSys.py` without any parameters, `recoverSys.py` performs a recovery to the most recent point in time.

i Note

Starting `recoverSys.py` on its own does not do anything.

Related Information

SAP HANA SQL and System Views Reference

10.2.5.2.4.1 Recover a Database Using Native SQL

Recovery using SQL statements is based on a full backup (complete data backup or data snapshot) in the backup catalog. You can specify a full backup or let SAP HANA decide which backups to recover from.

Procedure

To recover a system database (or an SAP HANA single-container system):

1. Set the environment using `HDBSettings.sh`.
2. Execute the `recoverSys.py` tool: `HDBSettings.sh recoverSys.py [<parameters>]`.

`recoverSys.py` shuts down the database.

Results

Once the master name server on the database has started, `recoverSys.py` terminates.

To check that the recovery was successful, see the `backup.log`.

For more information, see *Diagnosis Files for Backup and Recovery*.

i Note

If `recoverSys.py` returns an exit code '0', this is not confirmation that the recovery was successful.

The recovery is not complete yet. You still need to wait until the recovery has completed.

If you use the parameter `--wait`, the script waits until the recovery has completed.

If you do not use the `--wait` parameter, you need to manually check whether the recovery has completed by looking at the instance status or the logs.

Related Information

[Diagnosis Files for Backup and Recovery \[page 1272\]](#)

10.2.5.2.4.2 Options for Recovery with `recoverSys.py`

The default behavior of the `recoverSys.py` tool can be overridden using the options described below.

Options for `recoverSys.py`

<code>recoverSys.py</code> Options	Description
<code>--help</code>	Get help for the <code>recoverSys.py</code> script.

recoverSys.py Options

Description

`--command="<SQL command>"`

Use this option to specify a recovery command.

❖ Example

```
HDBSettings.sh recoverSys.py --  
command="RECOVER DATABASE UNTIL  
TIMESTAMP '2018-10-22 15:00:00'"
```

This statement performs a recovery to the database state of '2018-10-22 15:00:00'.

❖ Example

```
HDBSettings.sh recoverSys.py --  
command="RECOVER DATABASE UNTIL  
TIMESTAMP '2018-10-22 15:00:00' USING  
CATALOG PATH ('/remote/backup/CHH/  
catalog') USING BACKUP_ID  
1380740407446 CHECK ACCESS USING  
FILE"
```

This statement performs a recovery to the database state of '2018-10-22 15:00:00' based on the data backup identified by **BACKUP ID** '1380740407446', using the backup catalog located in '/remote/backup/CHH/catalog'.

The statement checks the availability of the backup files before actually performing the recovery.

❖ Example

To perform a recovery on a remote host, pass the recovery command to a remote shell command.

```
ssh <sid>adm@<remoteHost>  
"HDBSettings.sh recoverSys.py --  
command=\"RECOVER DATABASE UNTIL  
TIMESTAMP '2018-10-22 15:00:00'\""
```

→ Remember

The times specified are UTC times.

recoverSys.py Options	Description
--wait	<p>Causes the script to wait until the recovery has completed (either successfully or unsuccessfully).</p> <p>Default: The script does not wait for the recovery to complete. The recovery is started and runs in the background.</p> <p>If the script is terminated manually, the database recovery will not stop.</p> <p>For more information, see <i>Starting and Stopping Distributed SAP HANA Systems Using sapcontrol</i>.</p> <div data-bbox="826 701 1378 875" style="background-color: #f0f0f0; padding: 5px;"> <p>i Note</p> <p>If <code>recoverSys.py</code> is called automatically, you should use the option <code>--wait</code> to wait for the recovery to complete before you send further commands to the database.</p> </div>
--password=<password>	<p>If authentication is necessary, you can supply a password for <code><sid>adm</code>.</p> <p>If you do not specify the password, <code>recoverSys.py</code> prompts you to enter a password.</p> <div data-bbox="826 1113 1378 1234" style="background-color: #f0f0f0; padding: 5px;"> <p>i Note</p> <p>If you use the <code>--password</code> option, the password can be displayed in the process list of the operating system.</p> </div>
--timeout=<time>	<p>Specify a timeout for database shutdown and start.</p> <p>Default: 120s</p>
--licenseFile=<file name>	<p>Specify a license key file to append to the recovery command as a <code>SET LICENSE</code> clause.</p> <p>If you specify a command using the <code>--command</code> option, <code>SET LICENSE</code> is automatically appended to the command.</p> <div data-bbox="826 1612 1378 1854" style="background-color: #f0f0f0; padding: 5px;"> <p>i Note</p> <p>An SAP HANA license key becomes invalid if the <code><SID></code> or landscape ID is changed. The recovered system is assigned a temporary license that is valid for 90 days. You can apply to SAP to have the license from the source system transferred to a new license key for the recovered system.</p> </div>
--semaphoreOnly	<p>For use by SAP HANA cockpit and SAP HANA studio only.</p>

recoverSys.py Options	Description
<code>--masterOnly</code>	For use by SAP HANA cockpit and SAP HANA studio only.
<code>--forceMaster <host></code>	<p><code>recoverSys.py</code> attempts to use the current host as the master host for the recovery. If this host cannot be used as a master, <code>recoverSys.py</code> fails. To use a different host, use <code>--forceMaster</code> to specify the master host for the recovery.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>At most, three hosts can be used as the master host. The roles of the hosts are defined through the <code>name-server.ini</code> file. For this reason, it is not possible to use any random host as the master host.</p> </div>
<code>--feature</code>	For use by SAP HANA studio only.
<code>--silent</code>	Use this option to reduce diagnostics output.
<code>--cancel</code>	<p>Use this option to cancel a recovery after it has started.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>Canceling a recovery makes the database state inconsistent. SAP HANA prevents an inconsistent database from being started.</p> <p>To be able to work with the database again after canceling a recovery, you would need to perform the recovery again.</p> <p>For more information, see <i>Cancel a Recovery</i>.</p> </div>

Related Information

[Starting and Stopping Distributed SAP HANA Systems Using SAPControl \[page 1448\]](#)

[Cancel a Recovery \[page 1364\]](#)

10.2.5.2.5 Cancel a Recovery

You can cancel a recovery while it is in progress. After a recovery is canceled, it needs to be repeated or resumed before work can continue in the database.

Procedure

While a recovery is in progress, the option to cancel it is displayed.

1. Choose [Cancel Recovery](#) from the recovery progress view.
2. Confirm your decision.

i Note

It is possible to cancel a recovery from the command line using `recoverSys.py`.

For more information, see [Options for Recovery with recoverSys.py](#).

Results

⚠ Caution

After a recovery has been canceled, the database has an inconsistent state, and it will not be possible to start the database.

SAP HANA prevents a database with an inconsistent state from being started.

The only way to make the database available is to repeat or resume the recovery.

For more information about resuming a recovery, see [Resume a Canceled Recovery](#).

i Note

If you attempt to restart the database after a recovery is canceled, the following message is written to the nameserver trace file:

```
Cannot start the service 'nameserver' at '<host:SQL Port>' responsible for the volume '<volume number>' because of an error during recovery.
```

Related Information

[Options for Recovery with recoverSys.py \[page 1360\]](#)

[Resume a Canceled Recovery \[page 1365\]](#)

10.2.5.2.6 Resume a Canceled Recovery

Instead of repeating an entire recovery, it is possible to resume a recovery that was canceled or aborted. It is normally only necessary to resume a recovery in exceptional circumstances.

Prerequisites

i Note

After a recovery is canceled or aborted, the SAP HANA database cannot start again. Before work can continue in the database, the recovery must be completed.

! Restriction

A recovery from **only a full data backup** cannot be resumed.

If a recovery from only a full data backup is canceled or fails, the recovery needs to be repeated from the beginning.

- It is possible to resume a recovery that uses a full backup (a data backup or a data snapshot) with delta backups and log backups.
A recovery can only be resumed after the recovery from the full backup is completed.
If the recovery is canceled or fails during recovery from a data backup, the recovery cannot be resumed.
- To resume a recovery from a data snapshot, you can use the same data snapshot that you used when you started the recovery.
The data snapshot does not need to be replicated to the data area again. You only need to ensure that the required delta backups and log backups are available.
- It is possible to resume a recovery both with file-based backups and with third-party backup tools.

Disk Sizing

When you resume a recovery, disk space requirements can temporarily increase. If you resume a recovery, as a general rule, ensure that you are working with enough disk space to contain at least the data backup and the delta backups.

Context

In many situations, a recovery can be repeated from the beginning in only a short time, and the database can be running again with only minimal loss of uptime. However, in some situations, having to repeat a recovery from the beginning may cause a significant delay. In such situations, the option to resume a recovery can save a considerable amount of time, both with a very large database or a relatively small database.

Fallback Points

During a recovery, SAP HANA automatically defines fallback points, which mark the points after which it is possible to resume a recovery.

The first fallback point occurs after the recovery from the complete data backup. After recovery from a delta backup is completed, another fallback point is set. Fallback points are also written during log recovery.

A recovery can be resumed from the latest fallback point. The part of the recovery from before the fallback point does not need to be repeated.

i Note

If you resume a recovery to an earlier point in time than for the canceled recovery, the log fallback points are not used. In this situation, only the fallback point for the recovery from the data backup (and, if used, also the delta backups) can be used; the log recovery needs to be repeated from the beginning.

After a recovery has been successfully completed, the fallback points are invalidated. It is then no longer possible to perform a new recovery based on those fallback points.

Configuring Fallback Points

You can define how often fallback points are written during log recovery.

Configure the parameter `log_recovery_resume_point_interval` in the `global.ini` configuration file.

The time interval that you specify translates to the maximum acceptable database uptime lost while the log recovery is resumed after a recovery from the data backup (and possibly also the delta backups) has been completed.

Viewing Fallback Points

Each SAP HANA service defines its own service-specific fallback point after recovery from a data backup or delta backups. Together, these service-specific fallback points make up the fallback point for the whole database. When a fallback point has been written for **all the services** in the database, the fallback point for the whole database is recorded in `backup.log`.

i Note

No fallback points are written during log recovery.

❖ Example

The following example shows one fallback point in the `backup.log` file.

In this example, the fallback point is written to `fallback_databackup_0_1`:

```
2017-05-19T17:41:46+02:00 P012898 15c215ef3c7 INFO RECOVERY fallback point
written into /usr/sap/DB1/SYS/global/hdb/data/mnt00001/hdb00002.00004/
fallback_databackup_0_1
```

For more information, see `backup.log`.

Procedure

Resume a Canceled Recovery

1. Re-start the recovery.

You can use either SAP HANA cockpit or SAP HANA studio.

A partially completed recovery that can be resumed is indicated in the backup overview.

For more information, see *Recover a Database*.

2. In the recovery dialog, select the backup with which the recovery can be resumed.
3. Follow the steps described on-screen to complete the recovery.

Results

When the recovery is complete, a message confirms this, and shows the time to which the database was recovered.

The SAP HANA database is now online and can be used by applications.

Related Information

[backup.log \[page 1272\]](#)

[Recover a Database \[page 1347\]](#)

10.2.5.2.6.1 Resuming a Canceled Recovery Using Native SQL

You can resume a recovery by using the `recoverSys.py` tool and the SQL recovery command with `USING RESUME`.

For more information about `recoverSys.py`, see *Recovering a Database Using Native SQL*.

Note

`USING RESUME` can only be used if a fallback point already exists.

If no fallback point exists, an error is returned.

For simple SQL statements, `USING RESUME` is appended.

For SQL statements with options such as `IGNORE DELTA DATA BACKUPS` or `CHECK ACCESS`, `USING RESUME` should be placed after the path and before the option.

Sample Code

```
RECOVER DATABASE UNTIL TIMESTAMP '2017-05-19 17:41:46' USING RESUME
```

```
RECOVER DATABASE FOR Tenant_1 UNTIL TIMESTAMP '2017-05-19' USING CATALOG PATH  
( '/hana/DB1/backup/catalog' ) USING LOG PATH ( '/hana/DB1/backup/log' ) USING  
RESUME CHECK ACCESS ALL;
```

Related Information

[Recovering a System Database Using Native SQL \[page 1358\]](#)

10.2.5.3 Recovery Scenarios

Depending on the cause of the database failure, a different recovery strategy and procedure may be appropriate.

The following sections describe the recommended steps to recover the database in different recovery scenarios.

10.2.5.3.1 Data Area is Unusable (Disaster Recovery)

If the data area becomes unusable, you can recover an SAP HANA database.

If the data area is unusable, and all data changes since the last complete data backup are still available in the log backups and log area, you can still recover the data from committed transactions that was in the memory at the time of the failure. No committed data is lost.

⚠ Caution

If you reinstall the SAP HANA software to recover the database, **do not first create a data backup**.

The first data backup in a newly installed SAP HANA system creates a new backup catalog. As a result, the backup catalog from the old database is hidden and cannot be used for the recovery without manual intervention. However, in this scenario, the old backup catalog is still needed to recover the old database.

Once the database has been recovered successfully from a data backup or a data snapshot, the log entries from the log backups and the log area are replayed.

i Note

Currently, recovery of SAP HANA from a data snapshot is only supported for single-tenant systems.

It is also possible to recover the database using an older data backup or data snapshot in combination with delta backups and log backups. The log backups needed for the recovery include those created **after** the data backup or data snapshot.

For more information, see *SAP Note 1821207 (Determining required recovery files)*.

i Note

In the recovery dialog, ensure that the paths to the data and log backup files are correct.

Used for Recovery

- Data backup
Alternatively, data snapshot (for SAP HANA single-tenant systems only)
- Delta backups
- Log backups
- Log area

Steps for Recovery

Recover the database to its most recent state.

For more information, see *Recovering an SAP HANA Database*.

Related Information

[Recovering an SAP HANA Database \[page 1347\]](#)

[SAP Note 1821207](#)

10.2.5.3.2 Log Area is Unusable (Disaster Recovery)

If a log area becomes unusable, it is still possible to recover an SAP HANA database.

If the **log area** becomes unusable, it cannot be used for a recovery. It is only possible to recover the entries from the log backups. As a consequence, any changes that were made after the most recent log backup will be lost after a recovery. In addition, all the transactions that were open during the log backup will be rolled back.

It is still possible to recover the database to a point in time covered by the existing log backups.

Caution

If you reinstall the SAP HANA software to recover the database, **do not first create a data backup**.

The first data backup in a newly installed SAP HANA system creates a new backup catalog. As a result, the backup catalog from the old database is hidden and cannot be used for the recovery without manual intervention. However, in this scenario, the old backup catalog is still needed to recover the old database.

Note

Currently, recovery of SAP HANA from a data snapshot is only supported for single-tenant systems.

To prevent entries from the unusable log area from being replayed, in the recovery dialog, you must select the *Initialize log area* option. This option initializes the log area, and the old (unusable) content of the log area is lost.

Used for Recovery

- Data backup
Alternatively, data snapshot (for SAP HANA single-tenant systems only)
- Delta backups
- Log backups

Steps for Recovery

1. Recover the database to the most recent state.
When the database has been successfully recovered from the data backup or data snapshot, the log entries from the log backups are replayed.
2. Select the *Initialize log area* option.

Related Information

[Recovering an SAP HANA Database \[page 1347\]](#)

[Delta Backups \[page 1247\]](#)

10.2.5.3.3 Logical Error – Point-in-Time Recovery (Fault Recovery)

If a logical database error occurs, you can recover an SAP HANA database to a point in time before the error occurred.

i Note

Currently, recovery of SAP HANA from a data snapshots is only supported for single-tenant systems.

⚠ Caution

All changes made after the point in time of the recovery will be lost in the recovered database.

For this reason, a point-in-time recovery is not recommended for production systems.

If you need to perform a point-in-time recovery of your production system, consider recovering the database to a different system and importing the missing data back into your production system. For example, if a specific table was lost, import that table from the recovered system to the new system.

Used for Recovery

- Data backup from before the point in time to recover to.
Alternatively, data snapshot (for SAP HANA single-tenant systems only)
- Delta backups
- Log backups made after the data backup
- Log area

Steps for Recovery

Recover the database to a point in time before the logical error occurred.

i Note

You need to specify a point in time to which to recover the database. If you specify a point in time in the future, the effect is the same as recovering the database to the most recent state.

Related Information

[Recovering an SAP HANA Database \[page 1347\]](#)

10.2.5.3.4 Recovery with System Replication

If you are using a disaster-tolerant solution with system replication, some specific recovery scenarios apply.

Related Information

[Points to Note: System Replication \[page 1242\]](#)

10.2.5.3.4.1 Point-In-Time Recovery of a Primary System

A primary system in a system replication scenario can be recovered to a specific point in time.

To recover the primary system to a specific point in time (not to the most recent database state), you need to stop the secondary system for the time that the primary system is being recovered.

If the secondary system continues to run while the primary system is being recovered, the secondary system starts replication again immediately after the primary system is online again. As a consequence, incompatible, outdated log segments are sent to the secondary system.

To reinitialize system replication after the recovery, the offline secondary system must be registered again to the primary system, and restarted.

Used for Recovery

- Data backup
Alternatively, data snapshot
- Delta backups
- Log backups
The log backups that are associated with the data backup and cover the desired point-in-time (including the log backups made **after** the desired point in time).

Steps for Recovery

1. Stop the secondary system.
2. Recover the primary system.
3. Re-register the secondary system.
For more information, see *Configure the Secondary System*.
4. Start the secondary system.

Related Information

[Set Up SAP HANA System Replication with hdbnsutil \[page 1104\]](#)

[Recovering an SAP HANA Database \[page 1347\]](#)

[Recovery with System Replication \[page 1371\]](#)

10.2.5.3.4.2 Recovery of a New Primary System After a Takeover

With system replication, during a takeover, you switch your active system from the current primary system to the secondary system. After a takeover, it may become necessary to recover the active system (the former secondary system).

Used for Recovery

- Data backups
The data backup can be created either from the original primary or the now active system. Alternatively, you can use a data snapshot.

i Note

Currently, recovery of SAP HANA from a data snapshot is only supported for single-tenant systems.

To recover SAP HANA from a data snapshot, you need to recover the system database and the tenant database separately.

- Log backups
The log backups that belong to the data backup or the data snapshot. That is, if the data backup was made on the new primary system after takeover, only the log backups from the new primary system can be used.
- Log area of the new primary system

⚠ Caution

If SAP HANA is recovered from backups that were created with different UIDs, some third-party backup tools may prevent the recovery from being started.

For more information, contact your tool vendor or ensure that the same UID is used for all the backups used for a recovery.

Steps for Recovery

1. Ensure that the original primary system is stopped and is not writing complete data backups and log backups.
2. Ensure that the required data backup (or data snapshot) and the log backups can be accessed by the now active system.
3. Recover the now active system.
For more information, see *Recovering an SAP HANA Database*.
When the database has been successfully recovered from the data backup or data snapshot, the log entries from the log backups are replayed.

⚠ Caution

After a takeover, ensure that the original primary system does not continue to write log backups to the same location as the now active system.

For more information, see *Recovery with System Replication*.

i Note

After a takeover, it is not necessary to create a new full data backup of the now active system. Backups of the former primary system can be used to recover the database.

Related Information

[Recovering an SAP HANA Database \[page 1347\]](#)

[Recovery with System Replication \[page 1371\]](#)

10.2.6 Copying a Database Using Backup and Recovery

You can use backup and recovery to copy a system database or a tenant database within the same system or to a different system. A database copy is a quick way to set up a cloned database, for example, for training, testing, or development.

The following combinations of source database and target database can be used to create a database copy:

Source Database	Target Database
System database	The system database of a different system
Single-container system	Tenant database
Tenant database	A different tenant database in the same system A tenant database in a different system

i Note

An SAP HANA backup created with SAP HANA 1.0 SPS10 (single-container system) or newer can be used to recover a tenant database.

File System and Third-Party Backup Tools

You can copy an SAP HANA database using file-based backups or backups created using third-party tools.

i Note

To copy a database, it is not possible to mix backups from the different sources. The backup catalog, the data backups, and the log backups must be from either **only** a third-party backup tool or **only** the file system.

(To recover a database, it is possible to use a combination of backups from a third-party backup tool and backups from the file system, provided that the backups originate from the same SAP HANA database.)

Database Copy and Data Snapshots

Currently, it is only possible to use a data snapshot to back up and recover an SAP HANA single-tenant system.

To recover SAP HANA from a data snapshot, you first need to recover the system database, then the tenant database.

i Note

If you recover SAP HANA from a data snapshot, you must shut down the database **before** you make the data snapshot available in the data area of the storage system.

⚠ Caution

It is **not possible** to use a data snapshot of an SAP HANA single-container system to recover an SAP HANA multitenant database container.

Database Copy and System Replication

If you have system replication configured, and require near-zero downtime, consider using system replication to copy a tenant database.

For more information, see *Copying and Moving Tenant Databases Between Systems* in the *SAP HANA Administration Guide*.

Related Information

[Points to Note: Copying a Database Using Backup and Recovery \[page 1240\]](#)

[Copying and Moving Tenant Databases Between Systems \[page 1004\]](#)

[Copy a Database \[page 1378\]](#)

10.2.6.1 Prerequisites for Copying a Database Using Backup and Recovery

Before you can create a copy of an SAP HANA database, some important preparations are needed.

General Preparations for a Database Copy

- The version of the SAP HANA target database is the same or higher than the SAP HANA source database.
- You can copy a database to machines from different vendors and with different hardware configurations, provided that both the source and target machines are compliant with the SAP HANA appliance specifications.
Special requirements may apply to ensure the compatibility of SAP HANA backups with IBM Power Systems.
For more information, see *Points to Note: SAP HANA on IBM Power Systems*.
- To copy a complete SAP HANA system, the system database needs to be recovered first, and then all the tenant databases are recovered individually.
- For the system database, you must have the logon credentials of the operating system user (<sid>adm).
For a tenant database, the system database user must have the authorization DATABASE ADMIN.
- If you expect a different set of volumes to be recovered, before you start the recovery for a database copy, you should remove existing data and log volumes.
After a recovery to create a database copy, the system may include different volumes, or volumes may be assigned to different hosts.
Existing volumes that are not used for the new system will not be overwritten or removed. Any additional disk space is not released. This may lead to unexpected disk full situations.

Source Database

You can create a copy of a database using a complete data backup or a data snapshot. Additionally, using delta backups and log backups allows you to recover the database to a specific point in time.

- Make the data backups, the delta backups, and the log backup files from the source database available in the appropriate directory in the target database.
For more information, see *SAP Note 1821207 (Determining required recovery files)*.

i Note

If you wish to create a database copy using differential or incremental backups, you must also use log backups. If log backups are not available, you can only create a database copy using a full data backup.

- The content of the **log area** of the source database cannot be used for recovery.

⚠ Caution

With a database copy, the log area of the target system is always initialized. When the log area is initialized, the content of the log area is lost.

Target Database

- For a database copy using SAP HANA cockpit, the target database must be at least SAP HANA 2.0 SPS 01.
- The target database has sufficient disk space and memory.
The target system should have at least the same amount of disk space as the source system.
- The target system can have any number of hosts.
The number and type of services in tenant databases is set automatically during a copy from a complete data backup when a backup catalog is used.
If desired, and if performance limitations are acceptable, a copy of a database can be set up on a platform with less memory and CPU capacity and a different number of hosts.
- A valid license key is available for the target database.
For more information, see *Points to Note: License Key and Recovery*.

Database Copy Using Data Snapshots

⚠ Caution

Using a data snapshot, it is only possible to copy an SAP HANA system with a single tenant database.

If you attempt to use a data snapshot to copy an SAP HANA system with more than one tenant database, this may make the data area for all the tenant databases unusable.

- If you are using a data snapshot, you first need to stop the target SAP HANA database, then make available the data snapshot in the data area of the target database.
Currently, creating a copy of an SAP HANA from a data snapshot is only supported for single-tenant systems.
For more information, see *Data Snapshots*.
- For a database copy using data snapshots, the number of hosts and the number and type of services assigned to each host must be the same for the source database and the target database, and the mountpoint IDs must be identical.

For more information, see *Data Snapshots*.

Related Information

[Data Snapshots \[page 1249\]](#)

[Points to Note: License Key and Recovery \[page 1238\]](#)

[Points to Note: SAP HANA on IBM Power Systems \[page 1243\]](#)

SAP Notes

[SAP Note 1821207](#)

[SAP Note 1812980](#)

[SAP Note 2378962](#)

10.2.6.2 Copy a Database

Using SAP HANA cockpit, you can create a copy of an SAP HANA database by using backups of that database to recover to the same system or a different system.

In SAP HANA cockpit, follow the instructions on the screen. The actual sequence of steps that you perform depends on the specific options that you choose.

The options offered in SAP HANA cockpit are described in the sections below. To find out more about a set of options, expand its section.

Specify the Database to Copy To

System database	<p>In SAP HANA cockpit, go to the system database overview.</p> <p>Choose Copy Database.</p> <p>From here, follow the steps on the screen.</p> <p>You are prompted to shut down the system database in a later step.</p>
Tenant database	<p>In SAP HANA cockpit, go to Overall Tenant Statuses.</p> <p>Select the tenant database that you want to copy.</p> <p>Choose Copy Tenant.</p> <p>If the database is not already offline, you are prompted to shut it down.</p>

Specify the Database Copy Type

Full data backup only	<p>Create a copy of the database from the start time of the full data backup.</p>
Data and log backups	<p>Create a copy of the database to a specific point in time.</p> <p>In the next step, you are prompted to specify the time zone, the date and the time.</p>

Specify Whether to Use a Backup Catalog

If you are copying a database using a full data backup only, you can either select the data backup from the backup catalog, or specify its location without using a backup catalog.

A copy to a point in time is not possible if the full data backup is not recorded in the backup catalog.

No

Copy the database without using a backup catalog.

In the next steps, you are prompted to select the location of the data backup in the file system. The data backup to use is identified by its prefix.

File System:

❖ **Example**

For the tenant database PR1TENANT, the location of the data backup could look like this:

```
/usr/sap/PR1/HDB00/backup/data/  
DB_PR1TENANT/
```

Data Snapshot:

If you copy a database using a data snapshot, make the data snapshot available in the data area of the target system.

Backint:

Backups made using third-party tools always use the destination `/usr/sap/<SID>/SYS/global/hdb/backint`.

In the next steps, you are prompted to specify the source system and the backup prefix.

Yes

Use a backup catalog to locate the data backup.

In the next step, you are prompted to specify the location of the backup catalog.

File System:

❖ **Example**

For the tenant database PR1TENANT, if the backup catalog is in the same directory as the log backups, the location to specify could look like this:

```
/usr/sap/PR1/HDB00/backup/log/  
DB_PR1TENANT/
```

Backint:

Use the backup catalog in a third-party backup tool.

For more information, see *Working with Third-Party Backup Tools*.

Specify the Source System

Backint	If you copy a database using a third-party backup tool, you always need to specify a source system type.
System database	Specify the SID of the source system.
Tenant database	Specify whether to copy the tenant database from: <ul style="list-style-type: none">• A tenant database in the same SAP HANA system• A tenant database in a different SAP HANA system• A SAP HANA single-container system

i Note

For a database copy, it is **not possible** to mix backups from the different sources.

The backup catalog, the data backups, and the log backups must be from either **only** a third-party tool or **only** the file system.

(For a standard database recovery, it is possible to use a combination of backups from a third-party tool and the file system, provided that the backups originate from the same SAP HANA database.)

Specify the Backup to be Used

If you are using a backup catalog :	An overview of the available backups is displayed, depending on whether the catalog is in the file system or a third-party backup tool. It is not possible to mix backups from the different sources.
If you are not using a backup catalog :	Specify whether to use a data backup from the file system or a data snapshot. If you are using a data backup, in the next step, specify the location and the prefix of the backup. In the next step, if required, specify alternative locations for the data and log backups.

Check the Availability of the Backups

Optionally, you can check the availability of the backups to be used for the database copy.

This check ensures that the backups exist at the location specified.

Review Your Settings

Choose [Review](#) to display a summary of the settings you specified.

To display the SQL statement that will be used for the copy, choose [Display SQL Statement](#).

To change the settings, choose [Edit](#). All the settings that you specified are retained until you change them.

Start the Database Copy

If the settings are correct, choose [Start Copy](#).

SAP HANA cockpit displays a warning that you are about to overwrite the target system.

To start the database copy, choose [Start Copy](#) again.

The progress of the copy for each SAP HANA service is displayed.

While the database is being copied, it is possible to cancel the copy process.

Result

When the database copy is completed, a message confirms this.

A copy of the SAP HANA database is created in the location you specified.

If you copied a tenant database	The copy of the SAP HANA database is now online and can be used by applications.
If you copied a system database	You now need to copy the tenant databases in the SAP HANA system.

Note

For a database copy to a point in time, SAP HANA cockpit shows the point in time to which the copy was made.

The point in time that SAP HANA returns after a copy is the point in time of the last COMMIT to the database that has been copied.

For this reason, this point in time may be before the point in time that you specified for the copy.

Database Credentials

To allow SAP HANA cockpit to connect to the copied database, you may need to change the credentials of the user that is registered in SAP HANA cockpit.

Caution

Ensure that the correct passwords are used to connect to the copied SAP HANA database. If an incorrect password is used multiple times, SAP HANA may respond by locking that database user account.

Related Information

[Working with Third-Party Backup Tools \[page 1303\]](#)

[Cancel a Recovery \[page 1364\]](#)

10.2.6.3 Copy a Database to a Point in Time (SAP HANA Studio)

You can create a copy of a database using a complete data backup or a data snapshot. Additionally, using delta backups and log backups allows you to copy the database to a specific point in time.

Prerequisites

Before you can create a copy of an SAP HANA database, some important points must be considered.

For more information, see *Prerequisites for Copying a Database Using Backup and Recovery*.

Procedure

Start the recovery of the target database.

1. In SAP HANA studio, open the context menu and choose *Backup and Recovery*.

Then choose *Recover Tenant Database...*

Specify the recovery type.

2. You can specify either *Recover the database to its most recent state* or *Recover the database to the following point in time*.

If you are recovering the database to a specific point in time, enter the required information.

3. Choose *Next*.
4. If the log backups are not in the original location, specify a new location, and choose *Add*.

i Note

Do not change the SID of the *Source System* here. The SID of the source database is only relevant for database copy using third-party backup tools.

5. Choose *Next*.
6. Select the data backup or data snapshot.
7. Choose *Next*.
8. If the log backups are not in the original location, specify a new location, and choose *Add*.
9. Choose *Next*.

10. Select *Initialize log area*.
11. You can select *Install new license key* and specify the license key file or apply a license key file later.
12. Choose *Next*.
13. Review the recovery options and if correct, choose *Finish*.

The recovery is started.

Results

When the recovery has successfully completed, the database is started.

The source database has now been copied to the target database, and you can immediately begin work with the target database.

Caution

Ensure that the correct passwords are used to connect to the copied SAP HANA database. If an incorrect password is used multiple times, SAP HANA may respond by locking that database user account.

Related Information

[Prerequisites for Copying a Database Using Backup and Recovery \[page 1376\]](#)

[Steps After Copying a Database \[page 1389\]](#)

10.2.6.4 Copy a Database Using File-Based Data Backup or Data Snapshot Only (SAP HANA Studio)

You can create a copy of a database using only a data backup or a data snapshot. The content of the copied database is exactly the same as at the time at which the data backup or data snapshot was created.

Prerequisites

Before you can create a copy of an SAP HANA database, some important points must be considered.

For more information, see *Prerequisites for Copying a Database Using Backup and Recovery*.

Procedure

Start the recovery of the target database.

1. In SAP HANA studio, open the context menu and choose *Backup and Recovery*.

Then choose *Recover Tenant Database...*

2. Specify the recovery type *Recover the database to a specific data backup*.
3. Choose *Next*.
4. Specify the location of the backup catalog.

You can select *Recover using the backup catalog* and specify the location of the backup catalog, or select *Recover without the backup catalog*.

i Note

Do not change the SID of the *Source System* here. The SID of the source database is only relevant for database copy using third-party backup tools.

5. Choose *Next*.
6. Supply the required information.
Select the data backup or data snapshot, or select the *Destination Type File* and specify the location of the data backup and the backup prefix.
7. Choose *Next*.
8. Select *Initialize log area*.
9. You can select *Install new license key* and specify the license key file or apply a license key file later.
10. Choose *Next*.
11. Review the recovery options and if correct, choose *Finish*.

The recovery is started.

Results

When the recovery has successfully completed, the database is started.

The source database has now been recovered to the target database, and you can work with the recovered database.

Related Information

[Prerequisites for Copying a Database Using Backup and Recovery \[page 1376\]](#)

[Steps After Copying a Database \[page 1389\]](#)

10.2.6.5 Copying a Database Using Third-Party Backup Tools

Using third-party backup tools, you can create a homogeneous copy of an SAP HANA database.

Note

To create a copy of a database, it is not possible to mix backups from the file system and a third-party tool.

For a standard database recovery, it is possible to use a combination of backups from a third-party tool and the file system. The backups must originate from the same system.

Prerequisites

Before you can create a copy of a database, the following prerequisites must be met for the **target** database:

- The Backint agent is installed and configured.
For more information, consult the documentation from the tool vendor.
- If Backint parameter files are required, you need to:
 1. On operating system level, create the required Backint parameter files in locations that are accessible by the Backint agent.
One Backint parameter file is needed to access the backups from the source database.
A second Backint parameter file is needed to create new backups for the target database.

Example

If a source database SID is **ABC**, its parameter file for Backint could be called `param_backint_ABC.utl`. If a target database SID is **DEF**, its parameter file for Backint could be called `param_backint_DEF.utl`.

Note

For database copy, the SID must be included in the Backint parameter file name.

Make a note of the path and the parameter file name.

2. In SAP HANA studio, specify the path and name of the parameter file for data backups and, optionally, for log backups.
 1. To specify a parameter file in SAP HANA studio, while the target database is online, go to the [Backup Console](#) and choose **Configuration** > **Backint Settings** .
Specify the path and name of the parameter file as follows:

```
<path>/<user-defined>$(SAPSYSTEMNAME)<optional-extension>
```

Example

If the source database is called **ABC** and you have named the parameter file `param_backint_ABC.utl`, the path and name to specify here could be:

```
/usr/sap/DEF/SYS/global/hdb/opt/hdbconfig/param_backint_$(SAPSYSTEMNAME).utl
```

⚠ Caution

Ensure that the variable `$(SAPSYSTEMNAME)` appears exactly like this in the Backint parameter file name that you specify here. When a database copy using Backint is performed, `$(SAPSYSTEMNAME)` is dynamically replaced at runtime with the SID of the source database that you specify for database recovery.

To insert the name of the source tenant database into the file name, use the variable `$(DBNAME)`.

2. [Save](#) your changes.

- Move or delete available data backups and log backups from the **target** database in order to make them inaccessible during the recovery.
During a recovery, SAP HANA searches for the most recent backup catalog. If the backup catalog in the target database is newer than in the source database, SAP HANA may use an undesired backup for recovery.

10.2.6.5.1 Copy a Database to a Point in Time Using Third-Party Backup Tools

Using third-party backup tools with a data backup and log backups, you can create a copy of an SAP HANA database to a specific point in time.

Procedure

Start the recovery of the target database.

1. In SAP HANA studio, open the context menu and choose [Backup and Recovery](#).
2. Choose [Recover Tenant Database...](#)

Specify the recovery type.

3. You can specify either [Recover the database to its most recent state](#) or [Recover the database to the following point in time](#).

If you are recovering the database to a specific point in time, enter the required information.

4. Choose [Next](#).
5. Specify the SID of the [Source System](#).

The SID is used to copy an SAP HANA single-container database using backups created with third-party tools.

By default, the source is set to the SID of the target system. To specify the name of a tenant database, use `<DBNAME@SID>`.

i Note

The SID that you specify here is used by SAP HANA to replace the variable `<$(SAPSYSTEMNAME)>` in the [Backint Parameter File](#) displayed in the Backup Console.

To specify the name of the source tenant database, use the variable `<$ (DBNAME) >`.

6. Choose *Next*.

The available data backups for the specified source database are displayed.

7. Select a backup and choose *Next*.
8. For *Check Availability of Log Backups*, you can select *Third-Party Backup Tool (Backint)*.
9. Select *Initialize log area*.
10. You can select *Install new license key* and specify the license key file, or apply a license key file later.
11. Choose *Next*.
12. Review the recovery options and if correct, choose *Finish*.

The recovery is started.

Results

When the recovery has been successfully completed, the target database is started.

The source database has now been recovered to the target database, and you can work with the target database.

→ Tip

After the target database has been backed up, the source Backint parameter file is no longer needed. However, it is recommended that you retain the source Backint parameter file, as you will need it if you want to copy the source database again.

Related Information

[Steps After Copying a Database \[page 1389\]](#)

10.2.6.5.2 Copy a Database Using a Data Backup Only and Third-Party Backup Tools

Using third-party backup tools with a data backup only, you can create a homogeneous copy of an SAP HANA database.

Procedure

Start the recovery of the target database.

1. In SAP HANA studio, open the context menu and choose *Backup and Recovery*.
2. Choose *Recover Tenant Database...*
3. Specify *Recover the database to a specific data backup*.
4. Choose *Next*.
5. Specify the location of the backup catalog.

You can select *Recover using the backup catalog* and specify the location of the backup catalog, or select *Recover without the backup catalog*.

6. Specify *Backint System Copy* and specify the SID of the *Source System* that you want to copy.

To specify the name of a tenant database, use `<DBNAME@SID>`.

i Note

The SID that you specify here is used by SAP HANA studio to replace the variable `<$ (SAPSYSTEMNAME) >` in the *Backint Parameter File* displayed in the Backup Console.

To specify the name of the source tenant database, use the variable `<$ (DBNAME) >`.

7. Choose *Next*.
8. Supply the required information.
Select the data backup or data snapshot, or specify the location of the data backup and the backup prefix.
9. If the log backups are not in the original location, specify a new location, and choose *Add*.
10. Choose *Next*.
11. Select *Initialize log area*.
12. You can select *Install new license key* and specify the license key file, or apply a license key file later.
13. Choose *Next*.
14. Review the recovery options and if correct, choose *Finish*.

The recovery is started.

Results

When the recovery has been successfully completed, the target database is started.

The source database has now been recovered to the target database, and you can work with the target database.

i Note

After the target database has been backed up, the source Backint parameter file is no longer needed. However, if you need to copy the source database again, you will still need the source Backint parameter file.

Related Information

[Steps After Copying a Database \[page 1389\]](#)

10.2.6.6 Steps After Copying a Database

When you have completed a database copy, consider performing the following steps:

Storage

When the recovery is completed, it is recommended to back up the target database.

Nevertheless, it is recommended to keep the old backups available, at least until a new data backup of the target system has been created.

If you need to recover the target database before you have created a full backup of it, you can still use backups from the source database.

For this reason, immediately after the recovery, it is **not imperative** to create a new backup of the target database.

Services

An SAP HANA database automatically generates the services that it requires. You do not need to take any special steps to change the number of services.

After a database copy, you can remove a service. You may wish to do this, for example, if the target database has fewer hosts and more services than the source database.

i Note

If you use the SQL statement `ALTER DATABASE` to remove a service, the remaining services are not redistributed. After a service is removed, all the remaining services are where they were before.

For this reason, we strongly recommend that you use SAP HANA cockpit to remove a service.

For more information, see *Use the Cockpit to Add or Remove Services in a Tenant Database* in the *SAP HANA Administration Guide* and *ALTER DATABASE Statement (Tenant Database Management)* in the *SAP HANA SQL and System Views Reference*.

i Note

To remove a **host**, use SAP HANA database lifecycle manager (HDBLCM).

For more information about removing services, see *SAP HANA Administration - LCM Configuration* in the *SAP HANA Administration Guide*.

Backup Catalog

When a database copy is created, a new backup catalog is created in the target database. This new backup catalog allows the target database to be recovered using backups of the source database and new backups of the target database.

SAP HANA automatically uses the source database backups that are recorded in the backup catalog. You do not need to specify the SID of the source database again.

The backup catalog in the target database records **only** the backups from the source database that were used to recover the target database. If it is necessary to recover the target database again, you can only use the same backups of the source system that are recorded in the backup catalog. In this situation, older backups and different data backups, delta backups, and log backups cannot be used to recover the target database.

Scheduled Backups

After a database copy, ensure that any backups scheduled in the target database are configured in accordance with your requirements.

If backups were scheduled in the source database, after a database copy, the backups are scheduled to run in the target database with the same configuration as in the source database.

For more information, see *Schedule Backups*.

SAP HANA Secure User Store

If you use the SAP HANA secure user store (`hdbuserstore`) to connect to the database, you need to update the account information in the secure user store to match the accounts in the target database.

SAP HANA cockpit checks whether the user's logon credentials permit the user to perform a task.

For more information, see *Secure User Store (hdbuserstore)* in the *SAP HANA Administration Guide (Encryption)*.

Third-Party Backup Tools

After the target database has been backed up, the source Backint parameter file is no longer needed. However, it is recommended that you retain the source Backint parameter file, as you will need it if you want to copy the source database again.

Related Information

[Add or Remove Services in a Tenant Database \[page 255\]](#)

[Redistributing Tables in a Scaleout SAP HANA System \[page 600\]](#)

[Schedule Backups \[page 1325\]](#)

[Backup Catalog \[page 1255\]](#)

10.2.6.7 Database Copy: Scenarios

The following examples illustrate common situations in which a database copy is created.

10.2.6.7.1 Database Copy: Target Database has Fewer Hosts Than the Source Database

In this example scenario, we create a database copy where the target database has fewer hosts than the source database.

Prerequisites

The same prerequisites apply as for an SAP HANA database recovery.

For more information, see *Prerequisites for Database Recovery*.

Context

In the following example, the source database and the target database are both tenant databases.

The source database has two hosts, each with one index server.

The target database has only one host.

File-based backups are used.

Procedure

1. Create a target database with one host.

Follow the steps described in *Create a Tenant Database*.

2. Copy the database.

Follow the steps described in *Copy a Database Using File-Based Data Backup or Data Snapshot Only*.

All content of the data backup is recovered to the target host.

The target database now holds the data of the two hosts from the source database, and there are two index servers on the single-host target database.

3. Remove the index server.

For more information, see *Use the Cockpit to Add or Remove Services in a Tenant Database*.

Results

The source database with two hosts has been copied to the target database with one host.

Related Information

[Prerequisites for Database Recovery \[page 1332\]](#)

[Create a Tenant Database \[page 210\]](#)

[Copy a Database Using File-Based Data Backup or Data Snapshot Only \(SAP HANA Studio\) \[page 1383\]](#)

[Steps After Copying a Database \[page 1389\]](#)

[Add or Remove Services in a Tenant Database \[page 255\]](#)

10.2.6.7.2 Database Copy: Target Database has More Hosts Than the Source Database

In this example scenario, we use a third-party backup tool to create a database copy where the target database has more hosts than the source database.

Prerequisites

The same prerequisites apply as for an SAP HANA database recovery.

For more information, see *Prerequisites for Database Recovery*.

Context

In the following example, the source database and the target database are both tenant databases.

The source database has **two** hosts, each with one index server.

The target database has **three** hosts.

You are using a third-party backup tool.

Procedure

1. Create a target database with two hosts.

Follow the steps described in *Create a Tenant Database*.

At this stage, the third host in this constellation remains unused.

2. Copy the database.

Follow the steps described in *Copy a Database Using Data Backup Only and Third-Party Backup Tools*.

All content of the backup is recovered to the two existing hosts.

3. Add the remaining host to the target database.

For more information, see *Adding and Removing Hosts*.

4. Distribute all the data in the target database from the two hosts to the three hosts that are now available.

Results

A database with three hosts has been created, containing the data from the previous two-host database.

Related Information

[Prerequisites for Copying a Database Using Backup and Recovery \[page 1376\]](#)

[Create a Tenant Database \[page 210\]](#)

[Copy a Database Using a Data Backup Only and Third-Party Backup Tools \[page 1387\]](#)

[Steps After Copying a Database \[page 1389\]](#)

[Adding and Removing Hosts \[page 1406\]](#)

10.2.7 Planning Your Backup and Recovery Strategy

When you are planning a backup strategy, consider using a combination of data backups, automatic log backups, and data snapshots to minimize the risk of data loss, and to ensure that, if necessary, a recovery can be performed speedily.

When to Perform a Data Backup

It is recommended that you perform a data backup in the following situations:

- After the initial load
- At regular intervals

→ Tip

You can use less recent data backups for a recovery, provided that the subsequent log backups are available. If more log backups have to be replayed, the recovery takes longer to complete.

For this reason, we recommended that you use the most recent data backup and subsequent log backups to recover the database.

The more frequently a database is backed up, the faster the recovery will be.

- Before the database software is upgraded to a new version
If a software upgrade fails, it is possible to use the backup to recover the database to its state before the upgrade.

i Note

After an SAP HANA upgrade, the backup history is not broken. A full backup is not necessary to ensure that the backup history is intact.

- After any situation that causes log writing to be interrupted
For example, immediately after the log mode was changed.

Scheduling Regular Backups

It is strongly recommended to schedule regular data backups from the data area of your SAP HANA database to a secure location.

A possible backup scenario could look like this:

- Data snapshot: daily
You can create a data snapshot of an SAP HANA multitenant database container with one tenant database. For more information, see *Points to Note: Data Snapshots*.

i Note

If you have a backup and recovery strategy that is based on **data snapshots**, you must ensure that all data snapshots (or at least those you wish to use for a recovery) are replicated outside of the SAP HANA storage system.

- Data backup (file-based or with third-party tools): once a week
- Automatic log backups

For more information about scheduling backups using SAP HANA cockpit, see *Schedule Backups*.

Related Information

[Comparison of Data Backups and Data Snapshots \[page 1395\]](#)

[Points to Note About Backup and Recovery \[page 1231\]](#)

[SAP HANA Backup \[page 1245\]](#)

[SAP HANA Recovery \[page 1331\]](#)

[Points to Note: Data Snapshots \[page 1233\]](#)

[Schedule Backups \[page 1325\]](#)

10.2.7.1 Comparison of Data Backups and Data Snapshots

You can use this overview to help assess the benefits of using data backups and data snapshots as part of your backup strategy.

Comparison of Data Backups and Data Snapshots

	Data Backup to File	Data Backup Using Backint	Data Snapshot
Advantages	<ul style="list-style-type: none"> Integrity checks at block level For more information, see <i>Manually Checking Whether a Recovery is Possible</i>. Can be encrypted For more information, see <i>SAP HANA Backup Encryption</i>. 	<ul style="list-style-type: none"> Integrity checks at block level Integrated into existing data center infrastructure Backup tool offers additional features. For example, deduplication. Backups are immediately available for recovery. 	<ul style="list-style-type: none"> Fast Generates negligible network load Can be encrypted For more information, see <i>Encryption of Data Snapshots</i>.
Disadvantages	<ul style="list-style-type: none"> Requires additional storage Generates additional network load File system needs to be monitored (fill level) More time is needed to make backups available for recovery 	<ul style="list-style-type: none"> Generates additional network load 	<ul style="list-style-type: none"> No integrity checks at block level
Backup Size	<ul style="list-style-type: none"> Payload only 	<ul style="list-style-type: none"> Payload only 	<ul style="list-style-type: none"> Size of the data area (but is usually compressed or deduplicated by the storage tool)
Backup Duration	<ul style="list-style-type: none"> IO-bound (reading from data volume, writing to target) Network-bound (writing to target file system) 	<ul style="list-style-type: none"> IO-bound (reading from data volume) Network-bound (writing to backup server) 	<ul style="list-style-type: none"> Negligible (depending on the storage tool)

Related Information

[Manually Checking Whether a Recovery is Possible \[page 1335\]](#)

[SAP HANA Backup Encryption \[page 1252\]](#)

[Encryption of Data Snapshots \[page 1251\]](#)

10.2.8 Reference: Backup Console (SAP HANA Studio)

From the Backup Console in SAP HANA studio, you can display and configure backup settings.

Prerequisites

To work with the Backup Console, you need the system authorizations BACKUP ADMIN and CATALOG READ.

Open the Backup Console

1. In SAP HANA studio, go to the [Systems](#) view.
2. Open a system.
3. Double-click [Backup](#).
The Backup Console [Overview](#) is opened in a new tab.
If a backup is currently running, its status is displayed.

You can use the Backup Console to perform the following tasks:

Backup Console Tasks

Task	Steps
Start a full data backup or a delta backup	From the Overview tab, choose Open Backup Wizard .
Prepare a storage snapshot	From the Overview tab, choose Manage Storage Snapshot...
Monitor the progress of a running backup	<p>Go to the Overview tab.</p> <p>The status of data backups, storage snapshots, and delta backups is displayed. The status is displayed for backups started from SAP HANA studio as well as for backups started using SAP HANA cockpit or SQL statements.</p> <p>By default, the status is automatically refreshed every three seconds.</p> <p>To change this default refresh interval, in SAP HANA studio, choose ▶ Window ▶ Preferences ▶ SAP HANA ▶ Administration ▶ Backup Console ▶ Refresh Interval in Seconds ▾.</p> <p>Alternatively, you can refresh the overview manually by choosing Refresh in the toolbar.</p>
Review the most recent successful backup	Go to the Overview tab.

Task	Steps
Cancel a running backup	<p>Go to the Overview tab.</p> <p>If a backup is running, the option to cancel it is displayed.</p>
Display the content of the backup log	<p>From the Overview tab, choose Open Log File.</p>
Configure the backup settings	<p>Go to the Configuration tab.</p>
Backint settings.	<p>If you are using a third-party backup tool, you can specify a Backint parameter file for data backup and log backups.</p> <p>For more information, see Configure a Third-Party Backup Tool.</p>
Configure file-based data backup settings	<p>Go to File-Based Data Backup Settings.</p> <p>You can specify the default destination for file-based data backups.</p> <p>Each time you create a data backup, you have the option to use a different destination than the default.</p>

Task

Configure file-based log backup settings

Steps

Go to [Log Backup Settings](#).

Destination Type: Select either *File* or *Backint*.

i Note

The destination type *Backint* is only available if the Backint agent is installed.

Destination: Specify the default destination for log backups.

→ Tip

For improved data safety, it is recommended that you specify a path to a secure backup destination.

The data area, log area, data backups, and log backups should never be on the same physical storage devices.

Backup Interval: Specify the log backup interval.

Enable Automatic Log Backup: Disable or enable automatic log backups

⚠ Caution

During normal system operation (log mode `normal`), it is strongly recommended that you enable automatic log backups. When log segments are backed up, the space they occupied in the log area can be freed. SAP HANA can overwrite the newly freed space in the log area with new log entries. In this way, automatic log backups can prevent the log area from filling. If automatic log backups are disabled, the log area grows until the file system is full. If the file system is full, and no more log segments can be created, the database freezes.

Task

Set the maximum file size for file-based backup files

Steps

For file-based data backups, you may need to limit the maximum size of a single backup file. For example, due to file system limitations.

If the size of a data backup file for a service exceeds the specified limit, SAP HANA splits the file into multiple smaller files.

Go to the *Configuration* tab and select ► *File-Based Data Backup Settings* ► *Limit File Size* ►.

You can set the maximum file size in GB or TB.

The maximum file size applies to the data backups of all services.

The actual size of backups may be smaller than the specified maximum size.

i Note

If existing backups are overwritten by backups with the same names, at least **twice the space** in the backup location is needed, because the old backup and the new backup exist for a time in parallel.

→ Tip

The maximum file size is set by the `global.ini` parameter `data_backup_max_chunk_size`.

Task	Steps
Monitor the backup catalog	<p>Go to the Backup Catalog tab.</p> <p>The Backup Catalog tab displays an overview of backups. Here you can see the status of each catalog entry, together with its key information. To see the full details of an entry, select it in the list. The details are displayed in the Backup Details area. These include, for example, backup start and completion times, duration, size, throughput time, and a breakdown for each service.</p> <p>By default, only data backups and storage snapshots are displayed.</p> <p>To display delta backups, select Show Delta Backups.</p> <p>To display log backups, select Show Log Backups.</p> <p>The number of entries displayed is limited to 1000. You can change this setting in the Backup Console preferences (Window > Preferences > SAP HANA > Administration > Backup Console). Note that increasing the number of catalog entries displayed can impact the performance of the Backup Console.</p> <p>From the Backup Catalog tab, you can remove backup entries from the backup catalog, or also physically delete the backups.</p>

Related Information

- [SAP HANA Backup Types \[page 1246\]](#)
- [Backup Catalog \[page 1255\]](#)
- [Creating Backups \[page 1313\]](#)
- [Configure a Third-Party Backup Tool \[page 1304\]](#)

10.2.9 Reference: Backup Alerts

SAP HANA raises alerts that warn you of errors related to data and log backups.

Alert:	Check availability of volumes for backup
Alert ID:	34

Alert:	Check availability of volumes for backup
Description:	<p>This check warns you if a backup cannot be created because a volume or a service is unavailable.</p> <p>This alert can be triggered in combination with the alerts NOT_ASSIGNED_VOLUMES and CHECK_INACTIVE_SERVICES.</p>
Alert Text:	<no> (<service>) is not available. A backup cannot be created.
User Action:	Find out why the volume or service is not available.
Default Interval:	1 hour

Alert:	Check whether a data backup exists
Alert ID:	35
Description:	Checks whether at least one data backup exists, and warns you if no successful data backup is available for the instance. You are warned before any actual data loss occurs.
Alert Text:	No data backup exists.
User Action:	To ensure that your database can be recovered, create a data backup as soon as possible.
Default Interval:	6 hours
	This check is also performed when the database is started.

Alert:	Check last data backup
Alert ID:	36
Description:	<p>Checks whether the last data backup was successful, and warns you if the last data backup failed.</p> <p>If a scheduled backup fails, this check can help you prevent a situation from arising where no current backups are available.</p>
Alert Text:	The last data backup was not successful.
User Action:	Find out why the last data backup was not successful, resolve the problem, and create a new data backup as soon as possible.
Default Interval:	1 hour

Alert: Check the age of the last data backup	
Alert ID:	37
Description:	<p>Checks the age of the last successful data backup.</p> <p>If the last successful data backup is too old, the following alert levels are generated:</p> <ul style="list-style-type: none"> • 20 days: High • 7 days: Medium • 5 days: Low
Alert Text:	The last data backup is <days> days old.
User Action:	To reduce your downtime in a recovery situation, create a data backup as soon as possible.
Default Interval:	24 hours

Alert: Check last log backups	
Alert ID:	38
Description:	<p>Checks whether the last log backups were successful, and provides information about a failed log backup for a service or volume.</p> <p>As log backups are created automatically, this is the only way to notify users. This check should therefore be performed frequently and be accorded high priority.</p> <div style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>Log backups are created in separate backup destinations for each volume. For this reason, log backups need to be checked for each volume.</p> </div>
Alert Text:	The last log backup was not successful for volume <no> (<service> at <host>:<port>).
User Action:	Find out why a log backup was not successful and resolve the problem.
Default Interval:	15 minutes

Alert: Runtime of the currently running log backups	
Alert ID:	65
Description:	Determines whether the most recent log backup terminates in the specified time.

Alert: Runtime of the currently running log backups	
Alert Text:	A log backup with <ID <id> has been running for longer than <value> seconds>.
User Action:	Investigate why the log backup runs for too long, and resolve the issue.
Default Interval:	60 seconds

Alert: Storage snapshot (data snapshot) is prepared	
Alert ID:	66
Description:	Determines whether the period, during which the database is prepared for a data snapshot, exceeds a given threshold.
Alert Text:	The database was prepared for a data snapshot for longer than <value> seconds>.
User Action:	Investigate why the data snapshot was not confirmed or abandoned, and resolve the issue.
Default Interval:	5 minutes

Alert: Enable automatic log backup	
Alert ID:	69
Description:	Determines whether automatic log backup is enabled.
Alert Text:	Automatic log backup is disabled.
User Action:	Enable automatic log backup.
Default Interval:	15 minutes

Related Information

[Configure Backups \[page 1274\]](#)

[Estimate the Space Needed in the File System for a Data Backup \[page 1313\]](#)

[Data Snapshots \[page 1249\]](#)

[Monitoring Alerts \[page 373\]](#)

10.3 Scaling SAP HANA

There are two general approaches you can take to scale your SAP HANA system: scale up and scale out.

Scale up means increasing the size of one physical machine by increasing the amount of RAM available for processing.

Scale out means combining multiple independent computers into one system. The main reason for distributing a system across multiple hosts (that is, scaling out) is to overcome the hardware limitations of a single physical server. This allows an SAP HANA system to distribute the load between multiple servers. In a distributed system, each index server is usually assigned to its own host to achieve maximum performance. It is possible to assign different tables to different hosts (partitioning the database), as well as to split a single table between hosts (partitioning of tables).

Related Information

[Aspects of Scalability \[page 1404\]](#)

[Multiple-Host System Concepts \[page 1406\]](#)

[Host Addition Concepts \[page 1410\]](#)

[Adding Hosts to an SAP HANA System \[page 1413\]](#)

[Removing Hosts from an SAP HANA System \[page 1422\]](#)

[Configuring Host Roles \[page 1428\]](#)

[Configuring the Network for Multiple Hosts \[page 1438\]](#)

[Mapping Host Names for Database Client Access \[page 1076\]](#)

[Scaling SAP HANA Extended Application Services, Classic Model \[page 1448\]](#)

[Starting and Stopping Distributed SAP HANA Systems Using SAPControl \[page 1448\]](#)

10.3.1 Aspects of Scalability

Before you decide how to scale your SAP HANA implementation, there are a number of aspects that need to be considered, such as scaling data, performance, applications, and hardware.

Scaling the Data

One technique you can use to deal with planned data growth is to purchase more physical RAM than is initially required to set the allocation limit according to your needs, and then to increase it over time to adapt to your data. Once you have reached the physical limits of a single server, you can scale out over multiple machines to create a distributed SAP HANA system. You can do this by distributing different schemas and tables to different servers (complete data and user separation). However, this is not always possible, for example, when a single fact table is larger than the server's RAM size.

The most important strategy for scaling your data is **data partitioning**. Partitioning supports the creation of very large tables (billions of rows) by breaking them into smaller chunks that can be placed on different machines. Partitioning is transparent for most SQL queries and other data manipulations.

For more information, see the section on managing tables.

Scaling Performance

SAP HANA's performance is derived from its efficient, parallelized approach. The more computation cores your SAP HANA server has, the better the overall system performance is.

Scaling performance requires a more detailed understanding of your workload and performance expectations. Using simulations and estimations of your typical query workloads, you can determine the expected load that a typical SAP HANA installation may comfortably manage. At the workload level, a rough prediction of scalability can be established by measuring the average CPU utilization while the workload is running. For example, an average CPU utilization of 45% may indicate that the system can be loaded 2X before showing a significant reduction in individual query response time.

For more information, see the sections on workload management and performance analysis.

Scaling the Application

Partitioning can be used to scale the application as it supports an increasing number of concurrent sessions and complex analytical queries by spreading the calculations across multiple hosts. Particular care must be taken in distributing the data so that the majority of queries match partitioning pruning rules. This accomplishes two goals: directing different users to different hosts (load balancing) and avoiding the network overhead related to frequent data joins across hosts.

Scaling Hardware

SAP HANA is offered in a number of ways – in the form of an on-premise appliance, delivered in a number of different configurations and "sizes" by certified hardware partners or by using the tailored data center integration model, and as part of a cloud-based service. This creates different system design options with respect to scale-up and scale-out variations. To maximize performance and throughput, SAP recommends that you scale up as far as possible (acquire the configuration with the highest processor and memory specification for the application workload), before scaling out (for deployments with even greater data volume requirements).

i Note

The SAP HANA hardware partners have different building blocks for their scale-out implementations. Therefore, you should always consult with your hardware partner when planning your scale-out strategy.

Related Information

[Table Partitioning \[page 542\]](#)

[Workload Management \[page 621\]](#)

[Managing and Monitoring the Performance of SAP HANA \[page 394\]](#)

10.3.2 Adding and Removing Hosts

It is possible to add hosts after installation to a single-host or multiple-host SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM).

Before adding or removing hosts, it is important to review multiple-host system concepts, as well as the SAP HANA database lifecycle manager host addition concepts.

An SAP HANA system can also be configured as a multiple-host system during installation using the SAP HANA database lifecycle manager. For more information about installing an SAP HANA multiple-host system, see the *SAP HANA Server Installation and Update Guide*.

10.3.2.1 Multiple-Host System Concepts

It is important to review multiple-host system concepts like host grouping and storage options before installing a multiple-host system.

Host Types

When configuring a multiple-host system, the additional hosts must be defined as **worker** hosts or **standby** hosts (worker is default). Worker machines process data; standby machines do not handle any processing and instead just wait to take over processes in the case of worker machine failure.

Auto-Failover for High Availability

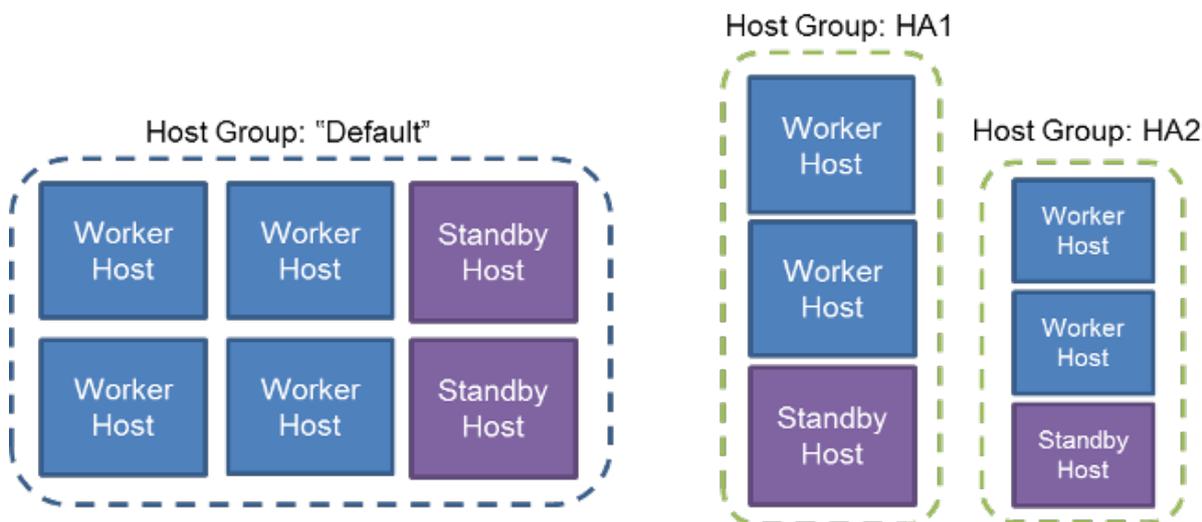
As an in-memory database, SAP HANA is not only concerned with maintaining the reliability of its data in the event of failures, but also with resuming operations with most of that data loaded back in memory as quickly as possible. Host auto-failover is a local fault recovery solution that can be used as a supplemental or alternative measure to system replication. One (or more) standby hosts are added to a SAP HANA system, and configured to work in standby mode.

Before installing a multiple-host system, it is important to consider whether high availability is necessary and how hosts should be grouped to ensure preferred host auto-failover. For host auto-failover to be successful, if the active (worker) host fails, the standby host takes over its role by starting its database instance using the persisted data and log files of the failed host. The name server of one of the SAP HANA instances acts as the

cluster manager that pings all hosts regularly. If a failing host is detected, the cluster manager ensures that the standby host takes over the role and the failing host is no longer allowed write access to the files (called fencing) so that they do not become corrupted. The crash of a single service does not trigger failover since services are normally restarted by `hdbdaemon`. For more information, see *Setting Up Host Auto-Failover* in the *SAP HANA Administration Guide*.

Host Grouping

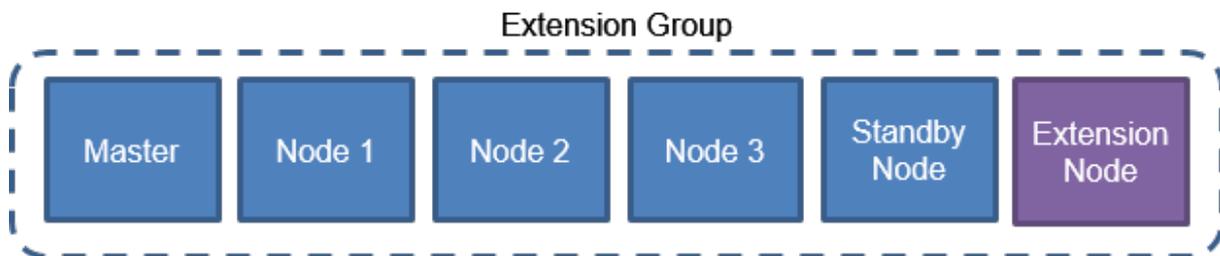
Host grouping does not affect the load distribution among worker hosts - the load is distributed among all workers in an SAP HANA system. If there are multiple standby hosts in a system, host grouping should be considered, because host grouping decides the allocation of standby resources if a worker machine fails. If no host group is specified, all hosts belong to one host group called "default". The more standby hosts in one host group, the more failover security.



If the standby hosts are each in a different host group, the standby host in the same group as the failing worker host is preferred. Only if no standby host is available in the same host group, the system will try to fail over to a standby host, which is part of another host group. The advantage of this configuration is that in an SAP HANA system with mixed machine resources, similar sized machines can be grouped together. If a small worker host fails, and a small standby in the same group takes over, the processes are moved to a machine with similar resources, which allows processing to continue as usual with optimal resource allocation.

Worker Host Grouping

If you use SAP Business Warehouse to apply a temperature-based data strategy you can significantly optimize the usage of memory and hardware resources by reserving one node of the scaled-out HANA landscape exclusively for warm data. Due to information lifecycle management, multi-temperature strategies are often applied, whereby data is classified by access frequency as either hot, warm or cold. Depending on this classification and data usage, this data is stored in different memory areas.



A multi-temperature memory strategy may be required for different reasons, for example:

- Storage of historical data
- Clickstream logs for multiple years of Web data and detailed machine logs
- Guidelines for saving company data, such as the need to save all data for at least seven years due to legal reasons

The standard SAP HANA sizing guidelines allow for a data footprint of 50% of the available RAM. This ensures that all data can be kept in RAM at all times and there is sufficient space for intermediate result sets. These sizing guidelines can be significantly relaxed on the extension group, since "warm" data is accessed

- less frequently,
- with reduced performance SLAs,
- with less CPU-intensive processes,
- only partially at the same time.

To implement a multi-temperature memory strategy, you can assign hosts to worker groups. Hot and warm data are then distributed across hosts. To increase performance and memory usage, a slave node is assigned to a separate "extension node". Unlike the standard nodes (Master and Slave), the extension node is intended exclusively for data that are not accessed as frequently (warm) as other data (hot). For more information, see *Data Temperature: Extension Node for Business Warehouse* in the *SAP HANA Administration Guide* and SAP Note 2453736.

Storage and File System Options

In single-host SAP HANA systems, it is possible to use local file systems residing on direct-attached internal or external storage devices, such as SCSI hard drives, SSDs, SAN storage, or NAS. However, in order to build a multiple-host system with failover capabilities this is not sufficient. Either the chosen file system type or the SAN Infrastructure along with a SAP HANA functionality capable of disc fencing must ensure the following:

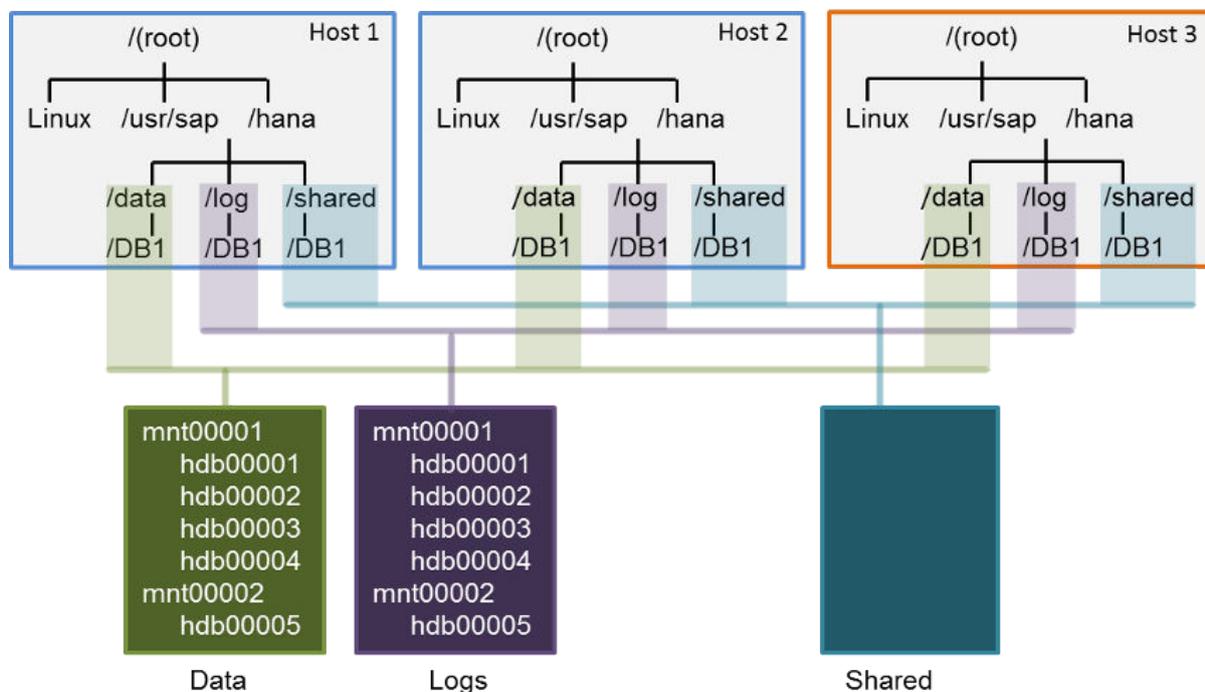
- The standby host has file access to data and log volumes of the failed host.
- The failed worker host no longer has access to write to files - called fencing.

There are two fundamentally different storage configurations which meet the two conditions above: **shared storage devices** or **separate storage devices with failover reassignment**. Do not confuse "shared storage" with the installation directory `/hana/shared` that must be shared across all hosts.

Shared File Systems

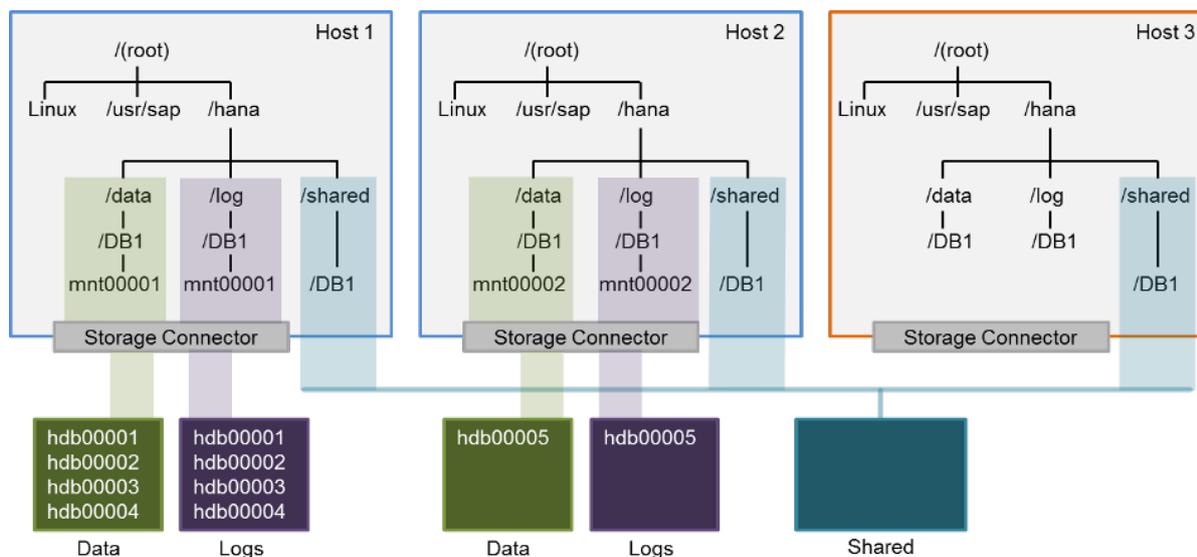
A shared storage subsystem, which is accessed using file systems such as NFS or IBM's GPFS, makes it easy to ensure that the standby host has access to all active host files in the system. In a shared storage solution, the externally attached storage subsystem devices are capable of providing dynamic mount points for hosts. Since shared storage subsystems vary in their handling of fencing, it is the responsibility of the hardware partner and their storage partners to develop a corruption-safe failover solution which is specific for the file system used to access that storage subsystem. An NFSv3 storage solution must be used in combination with the storage connector supplied by the hardware partner. NFSv4 and GPFS storage solutions can optionally be used with a storage connector.

A shared storage system could be configured as in the diagram below, however mounts may differ among hardware partners and their configurations. For more information, see the *SAP HANA Storage Whitepaper* available in SAP Note 1900823 in Related Information.



Non-shared Storage

It is also possible to assign every SAP HANA host a separate storage, which has nothing mounted except the shared area. A SAN storage must be used in combination with the SAP Fiber Channel Storage Connector, which SAP HANA offers storage technology vendors. During failover, SAP HANA uses the storage connector API to tell the storage device driver to re-mount the required data and logs volumes to the standby host and fence off the same volumes from the failed host.



In a non-shared environment, separate storage is used in combination with the storage connector API. For more information about the storage connector API, see the *SAP Fiber Channel Storage Connector Admin Guide* available in SAP Note 1900823 in Related Information.

Related Information

[SAP Note 405827](#)

[Setting Up Host Auto-Failover \[page 1208\]](#)

[SAP Note 1900823](#)

[Data Temperature: Extension Nodes \[page 581\]](#)

[SAP Note 2453736](#)

[More Details – HANA Extension Nodes for BW-on-HANA](#)

10.3.2.2 Host Addition Concepts

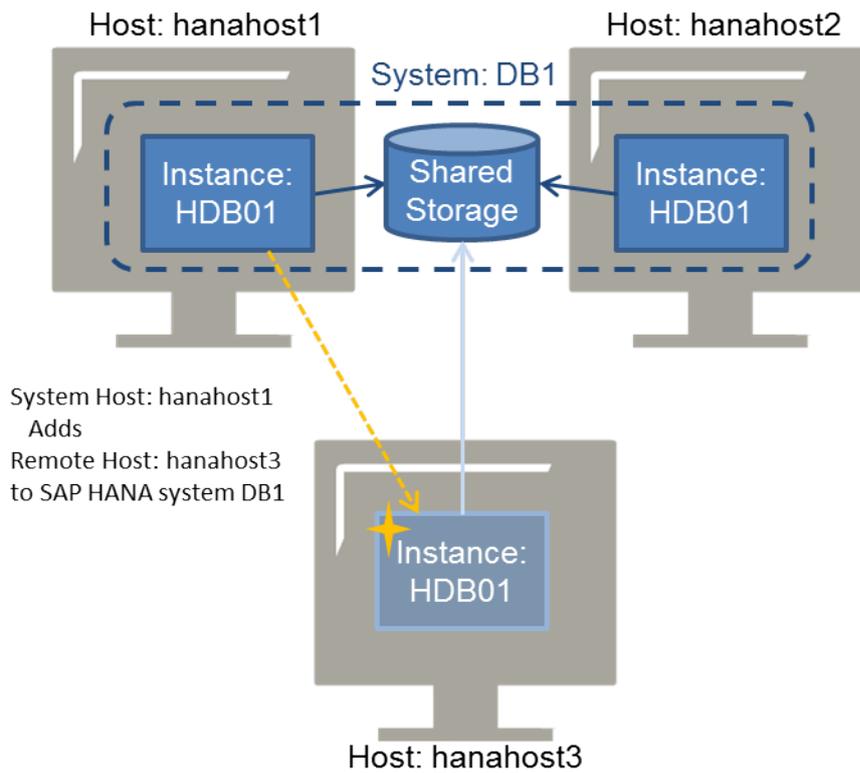
You can add hosts to an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM).

Using either the SAP HANA database lifecycle manager graphical user or command-line interface, one or multiple hosts can be added to an SAP HANA system in a variety of ways. The configuration options change depending on how the host is added.

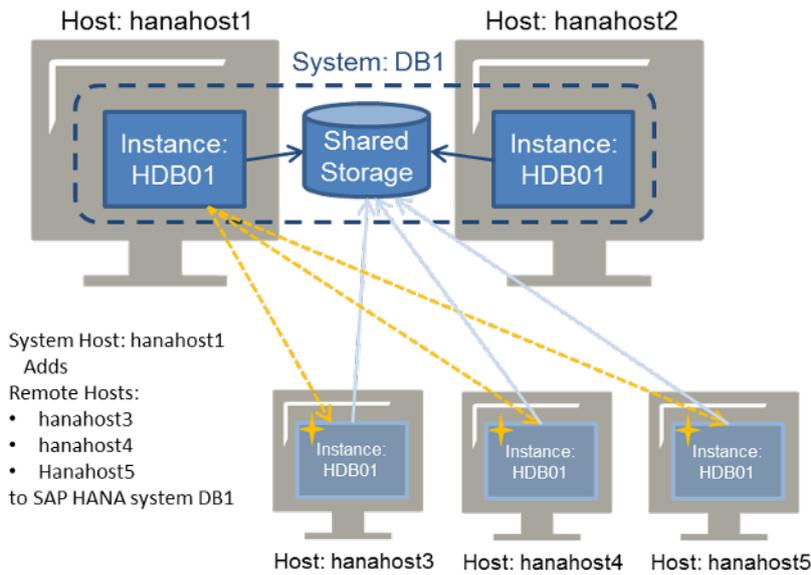
Adding Hosts from an Integrated Host

The first consideration is whether the host you are logged on to is integrated in the system. If you are logged on to a configured system host, then you are on an integrated host and adding a non-integrated host to the

system. In the diagram below, the hosts in the dotted line (hanahost1 and hanahost2) are integrated hosts because they both belong to the SAP HANA system DB1. Consider being logged on to hanahost1, and adding non-integrated host, hanahost3, to the SAP HANA system. The SAP HANA database lifecycle manager is started on the integrated host, hanahost1, and the addhost configuration task is carried out. The host information for hanahost3 is entered, and hanahost3 is configured as either a worker host or standby host. As soon as the addhost configuration task is finished, hanahost3 has access to the shared storage of the DB1 system.

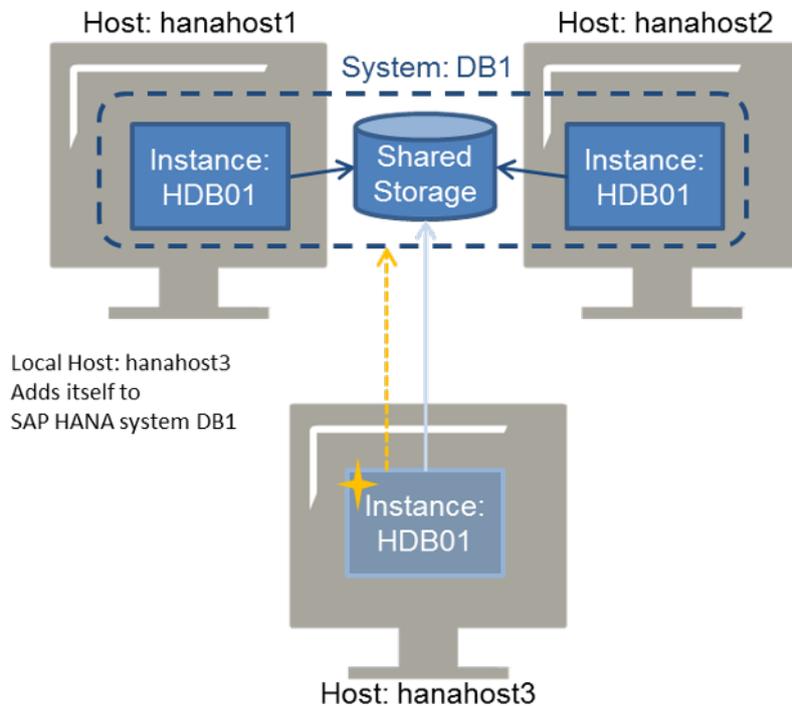


It is also possible to add multiple non-integrated hosts to the same system at one time. In the diagram below, three remote hosts (hanahost3, hanahost4, hanahost5) are added to the SAP HANA system (DB1) from a system host (hanahost1).



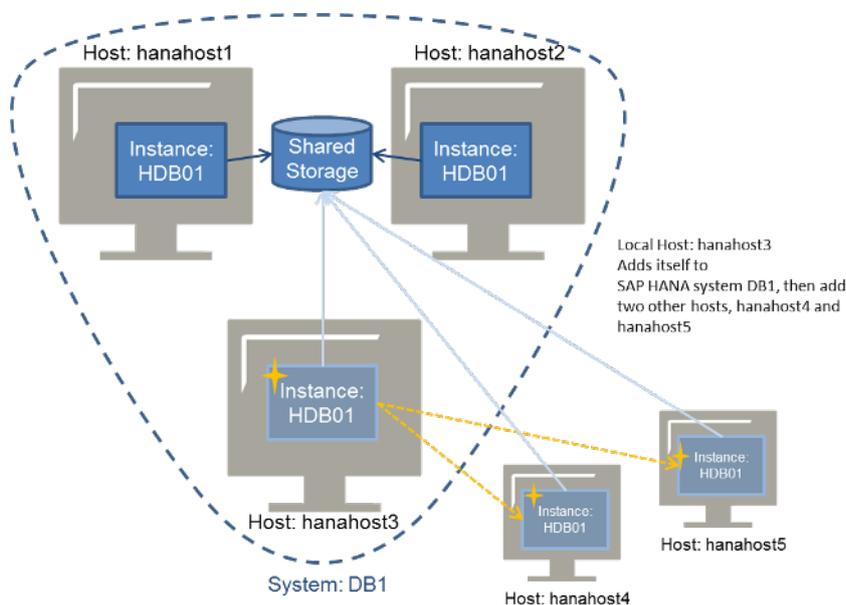
Adding Hosts from a Non-Integrated Host

Alternatively, a non-integrated host can add itself to an SAP HANA system. This is referred to as adding a host from a non-integrated host, because you are logged on to a host which you want to add to the system.



To add multiple hosts to an SAP HANA system from a non-integrated host, first the non-integrated host must be added (and, therefore, become integrated), and then it can add more hosts. The SAP HANA database lifecycle manager interface is designed so that the non-integrated host and the additional hosts can be added

in the same procedure. In the diagram below, the non-integrated host has already been newly added to the system (become integrated), and is now adding the other hosts.



Related Information

[Add Hosts Using the Command-Line Interface \[page 1416\]](#)

10.3.2.3 Adding Hosts to an SAP HANA System

You can add hosts to an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program or the SAP HANA database lifecycle manager (HDBLCM) Web user interface.

If you want to configure a new multiple-host (distributed) system during installation, see the multiple-host system installation information in the *SAP HANA Server Installation and Update Guide*.

Before adding a host to an SAP HANA system, you need to consider the following:

- If you are adding hosts from a host that is already integrated in the SAP HANA system
- If the system is a single-host or multiple-host system
- How many hosts you want to add to the system at one time

For more information about how these conditions affect the addition of hosts to an SAP HANA system see the host addition concepts in Related Information.

If you are adding a host to a single-host system, the listen interface is automatically configured to global during the host addition. After the host is added to the system, the internal network address can be defined and the inter-service communication can be reconfigured to a different setting, if required. For more information about configuring inter-service communication, see Related Information.

Related Information

[Multiple-Host System Concepts \[page 1406\]](#)

[Host Addition Concepts \[page 1410\]](#)

[Add Hosts Using the Graphical User Interface \[page 1414\]](#)

[Add Hosts Using the Command-Line Interface \[page 1416\]](#)

[Add Hosts Using the Web User Interface \[page 1419\]](#)

10.3.2.3.1 Add Hosts Using the Graphical User Interface

You can add hosts to an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in the graphical user interface.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).
- The host has access to the installation directories `<sapmnt>` and `<sapmnt>/<SID>`.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on as root user or as the system administrator user `<sid>adm`.
- The difference between the system time set on the installation host and the additional host is not greater than 180 seconds.
- The operating system administrator (`<sid>adm`) user may exist on the additional host. Make sure that you have the password of the existing `<sid>adm` user, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) resident program will not modify the properties of any existing user or group.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblcml
```

By default, `<sapmnt>` is `/hana/shared`.

2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

```
./hdblcmlgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

3. Select *Add Hosts to SAP HANA System* from the activity options. Then select *Next*.
4. Select *Add Host...* to define the required parameters. Then select *Next*.

Field Name	Description
<i>Host Name</i>	Specifies the host name of the machine.
<i>Role</i>	<p>Specifies the purpose of the SAP HANA host. Although multiple host roles may be assigned, check the corresponding documentation for the SAP HANA option for what configurations are supported in production environments.</p> <ul style="list-style-type: none"> ○ Database Worker (<i>worker</i>) - A worker host (default) is used for database processing. ○ Database Standby (<i>standby</i>) - A standby host is idle and available for fail-over in a high-availability environment. ○ Dynamic Tiering Worker (<i>extended_storage_worker</i>) - Worker host for SAP HANA dynamic tiering ○ Dynamic Tiering Standby (<i>extended_storage_standby</i>) - Standby host for SAP HANA dynamic tiering ○ Accelerator for SAP ASE Worker (<i>ets_worker</i>) - Worker host for SAP HANA accelerator for SAP ASE ○ Accelerator for SAP ASE Standby (<i>ets_standby</i>) - Standby host for SAP HANA accelerator for SAP ASE ○ Remote Data Sync (<i>rdsync</i>) - Host for SAP HANA remote data sync ○ Streaming Analytics (<i>streaming</i>) - Host for SAP HANA streaming analytics ○ XS advanced runtime worker (<i>xs_worker</i>) - Host for SAP HANA XS advanced runtime ○ XS advanced runtime standby (<i>xs_standby</i>) - Standby host for SAP HANA XS advanced runtime
<i>Worker Group</i>	Specifies the worker group of the host. If undefined, the worker group is named "default". If you are using extension node for Business Warehouse, you must name the worker group "worker_dt".
<i>High-Availability Group</i>	Specifies the host group ID for failover scenarios. If undefined, the host group is named "default".
<i>Storage Partition</i>	Specifies the storage partition number, which is a logical role number assigned to non-shared storage devices in a storage connector API. Standby hosts do not have a storage partition.

5. Define additional system properties.

Field Name	Description
<i>Inter-Service Communication</i>	<p>Specifies the listen interface for the internal network communication.</p> <p><i>global</i> - Binds the processes to all interfaces. This option does not require an internal network address entry.</p> <p><i>internal</i> - Binds the processes to this address only and to all local host interfaces. This option requires an internal network address entry.</p>
<i>Internal Network Address</i>	<p>Specifies the internal subset address in CIDR notation.</p> <p>If you define a value other than <i>local</i>, the local interfaces will always be open.</p>

Field Name	Description
<i>Certificate Host Name</i>	Specifies the hostname used for generation of self-signed SSL certificates for the SAP Host Agent.

- Review the summary, and select *Add Hosts* to finalize the configuration.

Results

You have added one or more new hosts to an SAP HANA system. The SAP HANA system you have configured is a multiple-host system.

The new hosts have been added to the SAP HANA landscape information. The new hosts have been added to the landscape information of the system database.

10.3.2.3.2 Add Hosts Using the Command-Line Interface

You can add hosts to an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in the command-line interface.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw,no_root_squash`).
- The host has access to the installation directories `<sapmnt>` and `<sapmnt>/<SID>`.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on as root user or as the system administrator user `<sid>adm`.
- The difference between the system time set on the installation host and the additional host is not greater than 180 seconds.
- The operating system administrator (`<sid>adm`) user may exist on the additional host. Make sure that you have the password of the existing `<sid>adm` user, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) resident program will not modify the properties of any existing user or group.

Procedure

- Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblc
```

By default, `<sapmnt>` is `/hana/shared`.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblcm --addhosts=<host>[,<host2>]
```

where the syntax for the `addhosts` call option is as follows:

```
<host name>:role=<role name>:group=<group ID>:storage_partition=<partition number>
```

Field Name	Description
<code><host name></code>	Specifies the host name of the machine.
<code>role</code>	Specifies the purpose of the SAP HANA host. Although multiple host roles may be assigned, check the corresponding documentation for the SAP HANA option for what configurations are supported in production environments. <ul style="list-style-type: none">◦ <code>worker</code> - A worker host (default) is used for database processing.◦ <code>standby</code> - A standby host is idle and available for failover in a high-availability environment.◦ <code>extended_storage_worker</code> - Worker host for SAP HANA dynamic tiering◦ <code>extended_storage_standby</code> - Standby host for SAP HANA dynamic tiering◦ <code>ets_worker</code> - Worker host for SAP HANA accelerator for SAP ASE◦ <code>ets_standby</code> - Standby host for SAP HANA accelerator for SAP ASE◦ <code>streaming</code> - Host for SAP HANA streaming analytics◦ <code>rdsync</code> - Host for SAP HANA remote data sync◦ <code>xs_worker</code> - Host for SAP HANA XS advanced runtime◦ <code>xs_standby</code> - Standby host for SAP HANA XS advanced runtime
<code>workergroup</code>	Specifies the worker group of the host. If undefined, the worker group is named "default". If you are using extension node for Business Warehouse, you must name the worker group "worker_dt".
<code>group</code>	Specifies the host group ID for failover scenarios. If undefined, the host group is named "default".
<code>storage_partition</code>	Specifies the storage partition number, which is a logical role number assigned to non-shared storage devices in a storage connector API. Standby hosts do not have a storage partition.

The required parameters depend on the type of host addition you are performing: host addition from an integrated host to a multiple-host system, host addition from an integrated host to a single-host system, or host addition from a non-integrated host. For more information about host addition types, see Related Information.

i Note

When using the command line, the options can be set interactively during configuration only if they are marked as interactive in the help description. All other options have to be specified in the command line. To call the help, in the `hdblcm` directory of the SAP HANA system, execute the following command:

```
./hdblcm --action=add_hosts --help
```

3. Select the index for the `add_hosts` action.
4. Define additional system properties.

Field Name	Description
<i>Inter-Service Communication</i>	Specifies the listen interface for the internal network communication. <code>global</code> - Binds the processes to all interfaces. This option does not require an internal network address entry. <code>internal</code> - Binds the processes to this address only and to all local host interfaces. This option requires an internal network address entry.
<i>Internal Network Address</i>	Specifies the internal subset address in CIDR notation. If you define a value other than <code>local</code> , the local interfaces will always be open.
<i>Certificate Host Name</i>	Specifies the hostname used for generation of self-signed SSL certificates for the SAP Host Agent.

5. Review the summary, and select `y` to finalize the configuration.

Results

You have added one or more new hosts to an SAP HANA system. The SAP HANA system you have configured is a multiple-host system.

The new hosts have been added to the SAP HANA landscape information. The new hosts have been added to the landscape information of the system database.

This configuration task can also be performed in batch mode and using a configuration file. For more information about the available configuration methods, see *Using the SAP HANA Platform LCM Tools*.

❁ Example

The following example adds two hosts, `Host1` and `Host2` to a single-host SAP HANA system. The role of the two hosts is `worker`, by default. No SSH keys are installed. A trusted connection between the hosts is configured and therefore, root user password is not required. The listen interface of the SAP HANA system is changed to `global`.

```
./hdblcm --action=add_hosts --addhosts=host1,host2 --root_user=lmroot --listen_interface=global
```

Related Information

[Host Addition Concepts \[page 1410\]](#)

[Configure SAP HANA Inter-Service Communication Using the Command-Line Interface \[page 1443\]](#)

[Using the SAP HANA Platform LCM Tools \[page 921\]](#)

[nstart \[page 198\]](#)

10.3.2.3.3 Add Hosts Using the Web User Interface

You can add hosts to an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) Web user interface.

Prerequisites

- On the host, which is to be added, SAP Host Agent is installed with SSL configured. The SAP Host Agent will create the `<sapsys>` group, if it does not exist prior to installation. Make sure that the group ID of the `<sapsys>` group is the same on all hosts.
- The difference between the system time set on the installation host and the additional host is not greater than 180 seconds.
- The operating system administrator (`<SID>adm`) user may exist on the additional host. Make sure that you have the password of the existing `<SID>adm` user, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) will not modify the properties of any existing user or group.
- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).
- The host has access to the installation directories `<sapmnt>` and `<sapmnt>/<SID>`.
- The SAP HANA system has been installed or updated with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
 - Microsoft Windows
 - Internet Explorer - Version 9 or higher
If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing [Tools > Compatibility View Settings](#) .
 - Microsoft Edge

- Mozilla Firefox - Latest version and Extended Support Release
- Google Chrome - Latest version
- SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
- Mac OS - Safari 5.1 or higher

i Note

For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: <code>https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html</code>

i Note

The URL is case sensitive. Make sure you enter upper and lower case letters correctly.

SAP HANA cockpit	<ol style="list-style-type: none"> 1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. <code>https://<host_FQDN>:<port></code>
-------------------------	---

i Note

FQDN = fully qualified domain name

2. Drill down on the name of the system from *My Resources* or from a group.
3. The links in *Platform Lifecycle Management* each launch additional functionality, giving you expanded capabilities for managing the resource.

2. Select the *Add Hosts* tile.
3. Optional: Modify the following parameters in the *Advanced Parameters Configuration* dialog. To access the *Advanced Parameters Configuration* dialog, click on the gear icon in the footer bar of the SAP HANA HDBLCM Web user interface.

Option	Description
import_xs_content	Imports SAP HANA XS advanced runtime content.
Install or Update SAP Host Agent	Installs or updates SAP Host Agent.
Do Not Start Added Hosts	Does not start hosts after addition.
Do Not Modify 'etc/sudoers' File	Prevents the file <code>/etc/sudoers</code> from being modified.
Timeouts	Sets customized timeouts (<code>start_instance</code> , <code>start_service</code>)

4. Provide the necessary credentials, then select *Add Host*.
5. Define the required host parameters. Then select *OK*.

Field Name	Description
<i>Host Name</i>	Specifies the host name of the machine.
<i>Role</i>	<p>Specifies the purpose of the SAP HANA host. Although multiple host roles may be assigned, check the corresponding documentation for the SAP HANA option for what configurations are supported in production environments.</p> <ul style="list-style-type: none"> ○ Database Worker (<i>worker</i>) - A worker host (default) is used for database processing. ○ Database Standby (<i>standby</i>) - A standby host is idle and available for fail-over in a high-availability environment. ○ Dynamic Tiering Worker (<i>extended_storage_worker</i>) - Worker host for SAP HANA dynamic tiering ○ Dynamic Tiering Standby (<i>extended_storage_standby</i>) - Standby host for SAP HANA dynamic tiering ○ Accelerator for SAP ASE Worker (<i>ets_worker</i>) - Worker host for SAP HANA accelerator for SAP ASE ○ Accelerator for SAP ASE Standby (<i>ets_standby</i>) - Standby host for SAP HANA accelerator for SAP ASE ○ Remote Data Sync (<i>rdsync</i>) - Host for SAP HANA remote data sync ○ Streaming Analytics (<i>streaming</i>) - Host for SAP HANA streaming analytics ○ XS advanced runtime worker (<i>xs_worker</i>) - Host for SAP HANA XS advanced runtime ○ XS advanced runtime standby (<i>xs_standby</i>) - Standby host for SAP HANA XS advanced runtime
<i>Worker Group</i>	Specifies the worker group of the host. If undefined, the worker group is named "default". If you are using extension node for Business Warehouse, you must name the worker group "worker_dt".
<i>High-Availability Group</i>	Specifies the host group ID for failover scenarios. If undefined, the host group is named "default".
<i>Storage Partition</i>	Specifies the storage partition number, which is a logical role number assigned to non-shared storage devices in a storage connector API. Standby hosts do not have a storage partition.

6. Define additional system properties.

Field Name	Description
<i>Inter-Service Communication</i>	<p>Specifies the listen interface for the internal network communication.</p> <p><i>global</i> - Binds the processes to all interfaces. This option does not require an internal network address entry.</p> <p><i>internal</i> - Binds the processes to this address only and to all local host interfaces. This option requires an internal network address entry.</p>

Field Name	Description
<i>Internal Network Address</i>	Specifies the internal subset address in CIDR notation. If you define a value other than <code>local</code> , the local interfaces will always be open.

7. Review the summary, and select *Run* to finalize the configuration.

Results

You have added one or more new hosts to an SAP HANA system. The SAP HANA system you have configured is a multiple-host system.

The new hosts have been added to the SAP HANA landscape information. The new hosts have been added to the landscape information of the system database.

Related Information

[SAPUI5 Developer Guide](#)

10.3.2.4 Removing Hosts from an SAP HANA System

You can remove hosts from an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program or the SAP HANA database lifecycle manager (HDBLCM) Web user interface.

Related Information

[Remove Hosts Using the Graphical User Interface \[page 1423\]](#)

[Remove Hosts Using the Command-Line Interface \[page 1424\]](#)

[Remove Hosts Using the Web User Interface \[page 1425\]](#)

10.3.2.4.1 Remove Hosts Using the Graphical User Interface

You can remove hosts from an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) in the graphical user interface.

Prerequisites

- You are logged in as root user.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- If you want to remove a host that runs the master name server, another host that will take over the role of the master name server must be up and running.
- You are logged on as root user or as the system administrator user `<sid>adm`.

⚠ Caution

Removing a host breaks the backup history of the database. To ensure that the database is fully recoverable, perform a full backup (data backup or storage snapshot) immediately after adding a service.

Procedure

1. Remove tenant-specific services. For more information, see *Remove a Service from a Tenant Database* in the *SAP HANA Tenant Databases Operations Guide*.
2. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblcmm
```

By default, `<sapmnt>` is `/hana/shared`.

3. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

```
./hdblcmmgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

4. Select *Remove Hosts from the SAP HANA System* from the activity options. Then select *Next*.
5. Select the host you would like to remove from the system. Then select *Next*.

You also have a choice to enable the following:

Field Name	Description
<i>Keep System Administrator User</i>	Keeps the system administrator user (<code><sid>adm</code>) from the source system to be used in the target system.
<i>Keep Home Directory of System Administrator</i>	Prevents the home directory of the source system administrator user (<code><sid>adm</code>) from being removed.

6. Enter the required credentials. Then select *Next*.

7. Review the summary, and select [Remove Hosts](#) to finalize the configuration.

Results

You have removed one or more new hosts from an SAP HANA system. This configuration task can also be performed using a configuration file. For more information about the available configuration methods, see *Using the SAP HANA Platform LCM Tools*.

Related Information

[Using the SAP HANA Platform LCM Tools \[page 921\]](#)

[Host Addition Concepts \[page 1410\]](#)

10.3.2.4.2 Remove Hosts Using the Command-Line Interface

You can remove hosts from an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) in the command-line interface.

Prerequisites

- You are logged in as root user.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- If you want to remove a host that runs the master name server, another host that will take over the role of the master name server must be up and running.
- You are logged on as root user or as the system administrator user `<sid>adm`.

⚠ Caution

Removing a host breaks the backup history of the database. To ensure that the database is fully recoverable, perform a full backup (data backup or storage snapshot) immediately after adding a service.

Procedure

1. Remove tenant-specific services. For more information, see *Remove a Service from a Tenant Database* in the *SAP HANA Tenant Databases Operations Guide*.
2. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdb1cm
```

By default, `<sapmnt>` is `/hana/shared`.

3. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdb1cm
```

4. Select the index for the `remove_hosts` action.
5. Select the hosts to be removed as a comma-separated list of indexes, and specify the following system properties:

Field Name	Description
<i>Keep System Administrator User</i>	Keeps the system administrator user (<code><sid>adm</code>) from the source system to be used in the target system.
<i>Keep Home Directory of System Administrator</i>	Prevents the home directory of the source system administrator user (<code><sid>adm</code>) from being removed.

6. Review the summary, and select `y` to finalize the configuration.

Results

You have removed one or more new hosts from an SAP HANA system. This configuration task can also be performed in batch mode and using a configuration file. For more information about the available configuration methods, see *Using the SAP HANA Platform LCM Tools*.

Related Information

[Using the SAP HANA Platform LCM Tools \[page 921\]](#)

[Host Addition Concepts \[page 1410\]](#)

10.3.2.4.3 Remove Hosts Using the Web User Interface

You can remove hosts from an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) Web user interface.

Prerequisites

- You are logged in as root user.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- If you want to remove a host that runs the master name server, another host that will take over the role of the master name server must be up and running.

- You are logged on as root user or as the system administrator user `<sid>adm`.

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
 - Microsoft Windows
 - Internet Explorer - Version 9 or higher
If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing **Tools > Compatibility View Settings**.
 - Microsoft Edge
 - Mozilla Firefox - Latest version and Extended Support Release
 - Google Chrome - Latest version
 - SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
 - Mac OS - Safari 5.1 or higher

Note

For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

Caution

Removing a host breaks the backup history of the database. To ensure that the database is fully recoverable, perform a full backup (data backup or storage snapshot) immediately after adding a service.

Procedure

1. Remove tenant-specific services. For more information, see *Remove a Service from a Tenant Database* in the *SAP HANA Tenant Databases Operations Guide*.
2. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: <code>https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html</code>

Note

The URL is case sensitive. Make sure you enter upper and lower case letters correctly.

Option	Description
SAP HANA cockpit	<ol style="list-style-type: none"> Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. <code>https://<host_FQDN>:<port></code> <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;"> <p>i Note FQDN = fully qualified domain name</p> </div> <ol style="list-style-type: none"> Drill down on the name of the system from <i>My Resources</i> or from a group. The links in <i>Platform Lifecycle Management</i> each launch additional functionality, giving you expanded capabilities for managing the resource.

- Select the *Remove Hosts* tile.
- Optional: Modify the following parameters in the *Advanced Parameters Configuration* dialog. To access the *Advanced Parameters Configuration* dialog, click on the gear icon in the footer bar of the SAP HANA HDBLCM Web user interface.

Option	Description
Do Not Remove XS Advanced OS Users	Prevents the XS advanced runtime OS Users from being removed.
Do Not Modify 'etc/sudoers' File	Prevents the file <code>/etc/sudoers</code> from being modified.
Timeouts	Sets customized timeouts (<code>start_instance</code> , <code>start_service</code> , <code>stop_instance</code> , <code>stop_service</code>).

- Select the host you would like to remove from the system. Then select *Next*.
You also have a choice to enable the following:

Field Name	Description
<i>Keep System Administrator User</i>	Keeps the system administrator user (<code><sid>adm</code>) from the source system to be used in the target system.
<i>Keep Home Directory of System Administrator</i>	Prevents the home directory of the source system administrator user (<code><sid>adm</code>) from being removed.

- Enter the relevant credentials. Then select *Next*.
- Review the summary, and select *Run* to finalize the configuration.

Related Information

[SAPUI5 Developer Guide](#)

[Using the SAP HANA Platform LCM Tools \[page 921\]](#)

[Host Addition Concepts \[page 1410\]](#)

10.3.3 Configuring Host Roles

It is possible to add and remove host roles after installation in a single-host or multiple-host SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM).

Before adding a host role, it is important to review multiple-host system concepts, and also to review the SAP HANA database lifecycle manager host addition concepts.

An SAP HANA system can also be configured with multiple host roles on single hosts during installation using the SAP HANA database lifecycle manager. For more information about installing an SAP HANA multiple-host system, see the *SAP HANA Server Installation and Update Guide*.

Related Information

[Adding Host Roles \[page 1428\]](#)

[Removing Host Roles \[page 1434\]](#)

10.3.3.1 Adding Host Roles

You can add host roles to hosts in an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program.

Related Information

[Add Host Roles Using the Graphical User Interface \[page 1428\]](#)

[Add Host Roles Using the Command-Line Interface \[page 1430\]](#)

[Add Host Roles Using the Web User Interface \[page 1431\]](#)

10.3.3.1.1 Add Host Roles Using the Graphical User Interface

You can add host roles to hosts in an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in the graphical user interface.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).

- The host has access to the installation directories `<sapmnt>` and `<sapmnt>/<SID>`.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on as root user or as the system administrator user `<sid>adm`.
- The difference between the system time set on the installation host and the additional host is not greater than 180 seconds.
- The operating system administrator (`<sid>adm`) user may exist on the additional host. Make sure that you have the password of the existing `<sid>adm` user, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) resident program will not modify the properties of any existing user or group.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblcml
```

By default, `<sapmnt>` is `/hana/shared`.

2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

```
./hdblcmlgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

3. Select [Add Host Roles](#) from the activity options. Then select [Next](#).
4. Select [Assign Roles...](#) to assign additional host roles to each host. Then select [Next](#).

Field Name	Description
<i>Role</i>	<p>Specifies the purpose of the SAP HANA host. Although multiple host roles may be assigned, check the corresponding documentation for the SAP HANA option for what configurations are supported in production environments.</p> <ul style="list-style-type: none"> ○ Database Worker (worker) - A worker host (default) is used for database processing. ○ Database Standby (standby) - A standby host is idle and available for fail-over in a high-availability environment. ○ Dynamic Tiering Worker (extended_storage_worker) - Worker host for SAP HANA dynamic tiering ○ Dynamic Tiering Standby (extended_storage_standby) - Standby host for SAP HANA dynamic tiering ○ Accelerator for SAP ASE Worker (ets_worker) - Worker host for SAP HANA accelerator for SAP ASE ○ Accelerator for SAP ASE Standby (ets_standby) - Standby host for SAP HANA accelerator for SAP ASE ○ Remote Data Sync (rdsync) - Host for SAP HANA remote data sync ○ Streaming Analytics (streaming) - Host for SAP HANA streaming analytics ○ XS advanced runtime worker (xs_worker) - Host for SAP HANA XS advanced runtime ○ XS advanced runtime standby (xs_standby) - Standby host for SAP HANA XS advanced runtime

5. Review the summary, and select *Run* to finalize the configuration.

10.3.3.1.2 Add Host Roles Using the Command-Line Interface

You can add host roles to hosts in an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in the command-line interface.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).
- The host has access to the installation directories `<sapmnt>` and `<sapmnt>/<SID>`.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on as root user or as the system administrator user `<sid>adm`.
- The difference between the system time set on the installation host and the additional host is not greater than 180 seconds.
- The operating system administrator (`<sid>adm`) user may exist on the additional host. Make sure that you have the password of the existing `<sid>adm` user, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) resident program will not modify the properties of any existing user or group.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdbclm
```

By default, <sapmnt> is /hana/shared.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdbclm --action=add_host_roles
```

3. Select the hosts to which you would like to assign additional roles.
4. Select the additional host roles that you want to assign for each host.

Field Name	Description
role	<p>Specifies the purpose of the SAP HANA host. Although multiple host roles may be assigned, check the corresponding documentation for the SAP HANA option for what configurations are supported in production environments.</p> <ul style="list-style-type: none">◦ <code>worker</code> - A worker host (default) is used for database processing.◦ <code>standby</code> - A standby host is idle and available for failover in a high-availability environment.◦ <code>extended_storage_worker</code> - Worker host for SAP HANA dynamic tiering◦ <code>extended_storage_standby</code> - Standby host for SAP HANA dynamic tiering◦ <code>ets_worker</code> - Worker host for SAP HANA accelerator for SAP ASE◦ <code>ets_standby</code> - Standby host for SAP HANA accelerator for SAP ASE◦ <code>streaming</code> - Host for SAP HANA streaming analytics◦ <code>rdsync</code> - Host for SAP HANA remote data sync◦ <code>xs_worker</code> - Host for SAP HANA XS advanced runtime◦ <code>xs_standby</code> - Standby host for SAP HANA XS advanced runtime

5. Enter the required credentials.
6. Review the summary, and select `y` to finalize the configuration.

10.3.3.1.3 Add Host Roles Using the Web User Interface

You can add host roles to hosts in an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in the Web user interface.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).

- The host has access to the installation directories `<sapmnt>` and `<sapmnt>/<SID>`.
- The SAP HANA system has been installed or updated with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
 - Microsoft Windows
 - Internet Explorer - Version 9 or higher
If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing **Tools > Compatibility View Settings**.
 - Microsoft Edge
 - Mozilla Firefox - Latest version and Extended Support Release
 - Google Chrome - Latest version
 - SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
 - Mac OS - Safari 5.1 or higher

i Note

For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: <code>https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html</code>
	<div data-bbox="411 1771 1401 1872" data-label="Text"> <p>i Note The URL is case sensitive. Make sure you enter upper and lower case letters correctly.</p> </div>
SAP HANA cockpit	<ol style="list-style-type: none"> 1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. <code>https://<host_FQDN>:<port></code>

Option	Description
	<p>i Note</p> <p>FQDN = fully qualified domain name</p> <ol style="list-style-type: none"> Drill down on the name of the system from My Resources or from a group. The links in Platform Lifecycle Management each launch additional functionality, giving you expanded capabilities for managing the resource.

- Select the [Add Host Roles](#) tile.
- Optional: Modify the following parameters in the [Advanced Parameters Configuration](#) dialog. To access the [Advanced Parameters Configuration](#) dialog, click on the gear icon in the footer bar of the SAP HANA HDBLCM Web user interface.

Option	Description
Do Not Start Hosts After Addition of Roles	Does not start hosts after addition of roles.
Do Not Modify 'etc/sudoers' File	Prevents the file <code>/etc/sudoers</code> from being modified.
Timeouts	Sets customized timeouts (<code>start_instance</code> , <code>start_service</code> , <code>stop_instance</code> , <code>stop_service</code>).

- Select the hosts to which you would like to assign additional roles.
- Select the additional host roles that you want to assign for each host.

Field Name	Description
<code>role</code>	<p>Specifies the purpose of the SAP HANA host. Although multiple host roles may be assigned, check the corresponding documentation for the SAP HANA option for what configurations are supported in production environments.</p> <ul style="list-style-type: none"> <code>worker</code> - A worker host (default) is used for database processing. <code>standby</code> - A standby host is idle and available for failover in a high-availability environment. <code>extended_storage_worker</code> - Worker host for SAP HANA dynamic tiering <code>extended_storage_standby</code> - Standby host for SAP HANA dynamic tiering <code>ets_worker</code> - Worker host for SAP HANA accelerator for SAP ASE <code>ets_standby</code> - Standby host for SAP HANA accelerator for SAP ASE <code>streaming</code> - Host for SAP HANA streaming analytics <code>rdsync</code> - Host for SAP HANA remote data sync <code>xs_worker</code> - Host for SAP HANA XS advanced runtime <code>xs_standby</code> - Standby host for SAP HANA XS advanced runtime

- Enter the required credentials.
- Review the summary, and select [Add Roles](#) to finalize the configuration.

Related Information

[SAPUI5 Developer Guide](#)

10.3.3.2 Removing Host Roles

You can remove host roles from hosts in an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program.

Related Information

[Remove Host Roles Using the Graphical User Interface \[page 1434\]](#)

[Remove Host Roles Using the Command-Line Interface \[page 1435\]](#)

[Remove Host Roles Using the Web User Interface \[page 1436\]](#)

10.3.3.2.1 Remove Host Roles Using the Graphical User Interface

You can remove host roles from hosts in an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in the graphical user interface.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).
- The host has access to the installation directories `<sapmnt>` and `<sapmnt>/<SID>`.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on as root user or as the system administrator user `<sid>adm`.
- The difference between the system time set on the installation host and the additional host is not greater than 180 seconds.
- The operating system administrator (`<sid>adm`) user may exist on the additional host. Make sure that you have the password of the existing `<sid>adm` user, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) resident program will not modify the properties of any existing user or group.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblcmlcm
```

By default, <sapmnt> is /hana/shared.

2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

```
./hdblcmlcmgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

3. Select *Remove Host Roles* from the activity options. Then select *Next*.
4. Select *Remove Roles...* to remove host roles from a host. Then select *Next*.
5. Review the summary, and select *Run* to finalize the configuration.

10.3.3.2 Remove Host Roles Using the Command-Line Interface

You can remove host roles from hosts in an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in the command-line interface.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw, no_root_squash`).
- The host has access to the installation directories <sapmnt> and <sapmnt>/<SID>.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on as root user or as the system administrator user <sid>adm.
- The difference between the system time set on the installation host and the additional host is not greater than 180 seconds.
- The operating system administrator (<sid>adm) user may exist on the additional host. Make sure that you have the password of the existing <sid>adm user, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) resident program will not modify the properties of any existing user or group.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdbclm
```

By default, <sapmnt> is /hana/shared.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdbclm --action=remove_host_roles
```

3. Select the hosts for which you would like to remove roles.
4. Select the host roles that you want to remove for each host.
5. Enter the required credentials.
6. Review the summary, and select **y** to finalize the configuration.

10.3.3.2.3 Remove Host Roles Using the Web User Interface

You can remove host roles from hosts in an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in the Web user interface.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw,no_root_squash`).
- The host has access to the installation directories <sapmnt> and <sapmnt>/<SID>.
- The SAP HANA system has been installed or updated with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
 - Microsoft Windows
 - Internet Explorer - Version 9 or higher
If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing **Tools > Compatibility View Settings**.
 - Microsoft Edge
 - Mozilla Firefox - Latest version and Extended Support Release

- Google Chrome - Latest version
- SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
- Mac OS - Safari 5.1 or higher

i Note

For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: <code>https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html</code>

i Note

The URL is case sensitive. Make sure you enter upper and lower case letters correctly.

SAP HANA cockpit	1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. <code>https://<host_FQDN>:<port></code>
-------------------------	--

i Note

FQDN = fully qualified domain name

2. Drill down on the name of the system from *My Resources* or from a group.
3. The links in *Platform Lifecycle Management* each launch additional functionality, giving you expanded capabilities for managing the resource.

2. Select the *Remove Host Roles* tile.
3. Optional: Modify the following parameters in the *Advanced Parameters Configuration* dialog. To access the *Advanced Parameters Configuration* dialog, click on the gear icon in the footer bar of the SAP HANA HDBLCM Web user interface.

Option	Description
Do Not Remove XS Advanced OS Users	Prevents the XS advanced runtime OS Users from being removed.
Do Not Start Hosts After Removal of Roles	Does not start hosts after removal of roles.
Do Not Modify 'etc/sudoers' File	Prevents the file <code>/etc/sudoers</code> from being modified.
Timeouts	Sets customized timeouts (<code>start_instance</code> , <code>start_service</code> , <code>stop_instance</code> , <code>stop_service</code>).

4. Select the hosts for which you would like to remove roles.
5. Select the host roles that you want to remove for each host. Then select *Next*.
6. Enter the relevant credentials. Then select *Next*.
7. Review the summary, and select *Remove Roles* to finalize the configuration.

Related Information

[SAPUI5 Developer Guide](#)

10.3.4 Configuring the Network for Multiple Hosts

As part of setting up a distributed system you should configure the network parameters to optimize performance. Make sure you do this before you add additional hosts because one server needs to be available so that you can connect to the SAP HANA studio.

You map host names to IP addresses by editing the section `internal_hostname_resolution` in the `global.ini` file.

General Network Layout



The figure shows a sample cluster with external addresses (10.68.22.*) and internal (192.168.2.*) addresses. To redirect the internal communication over the local network backbone, you could map the internal addresses to the host names of SAP HANA servers as shown in this example:

```
[communication]
listeninterface = .internal
#listeninterface = .global
#listeninterface = .local
#listeninterface = 192.168.2.0/24
[internal_hostname_resolution]
192.168.2.101 = hana01
192.168.2.102 = hana02
```

```
192.168.2.103 = hana03
```

For increased security, you can limit the binding of the processes in the `communication` section of the `global.ini` file. The option `listeninterface` can be set in one of the following ways:

- You can set it to one of the predefined keywords:
 - `.global`
 - `.internal`
 - `.local`

i Note

You must include the dot at the beginning of the keyword.

- You can set it to a subnet in CIDR notation (classless inter-domain routing).

The `.global` keyword (default) lets the process bind to all interfaces. The `.local` keyword opens the communication ports for internal usage on the local interfaces (which are 127.0.0.1 in IPv4 notation). This configuration is only an option for single host installations as the server is not reachable from the outside. These two options do not require a valid `internal_hostname_resolution` section.

If you specify a keyword other than `.local`, or if you specify a list of networks in CIDR notation, the local interfaces will always be open.

With the `.internal` setup, an `internal_hostname_resolution` section is required. This configuration scans `internal_hostname_resolution` for the local address of the host. The process is bound to this address only (and to all localhost interfaces). So you should add all hosts and their respective addresses to the `global.ini` immediately after installation of the first server. The SAP HANA instance on the first server then needs to be restarted for the changes to take effect. After that, the remaining hosts may be added.

With this configuration the whole landscape uses the internal network immediately after installation. To reduce the possibility of errors, it is also possible to install the whole landscape first without SAP HANA network configuration. This lets you run tests first before you establish the network. Then the configuration options remain the same and the whole SAP HANA landscape needs to be restarted for your changes to take effect.

It is possible to monitor the network using the [Monitor Network](#) link in the SAP HANA cockpit. The [Measure Network Speed](#) link on the [Monitor Network](#) page offers the possibility to measure the network speed between the hosts in a scale-out SAP HANA database. The [Network Speed Check](#) list offers an overview of all network channels between the involved hosts starting with the slowest network connection.

Related Information

[Internal Host Name Resolution \[page 1072\]](#)

10.3.4.1 Configuring SAP HANA Inter-Service Communication

In addition to external network connections, SAP HANA uses separate, dedicated connections exclusively for internal communication. These internal communication channels can be defined using the SAP HANA database lifecycle manager.

In a multiple-host system environment, inter-service communication takes place between the hosts of a multiple-host system on one site. Certified SAP HANA hosts contain a separate network interface card that is configured as part of a private network, using separate IP addresses and ports.

To prevent unauthorized access to the database via the internal communication channels in multiple-host systems, you can isolate internal network ports from client network. To do so, you route communication between the hosts of a multiple-host environment onto a specified network and bind those internal network services exclusively to the network interface.

In addition, this feature can now be used in the presence of a secondary site (system replication scenario). However, note that additional ports used for communication between primary and secondary sites are opened on the network interface. These ports need to be protected.

i Note

In single-host scenarios, the same communication channels are used for communication between the different processes on a single host. The internal IP addresses/ports are by default bound to the local interface. In multi-host scenarios, the specified network prefix must point to a network shared by all hosts. For security reasons, the network should belong to an internal network.

Related Information

[Configure SAP HANA Inter-Service Communication Using the Graphical User Interface \[page 1441\]](#)

[Configure SAP HANA Inter-Service Communication Using the Command-Line Interface \[page 1443\]](#)

[Configure SAP HANA Inter-Service Communication Using the Web User Interface \[page 1445\]](#)

10.3.4.1.1 Configure SAP HANA Inter-Service Communication Using the Graphical User Interface

To prevent unauthorized access to the SAP HANA system via the internal communication channels in multiple-host systems, you can configure inter-service communication using the SAP HANA database lifecycle manager graphical user interface.

Prerequisites

- The SAP HANA system has been installed or updated with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on with the required root user or system administrator user `<sid>adm` credentials.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblc
```

By default, `<sapmnt>` is `/hana/shared`.

2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

```
./hdblcgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

3. Select *Configure Inter-Service Communication* from the activity options. Then select *Next*.
4. Define the required parameters. Then select *Next*.

Field Name	Description
<i>Inter-Service Communication</i>	<p>Specifies the listen interface for the internal network communication between the services of an SAP HANA system. This is not related to the communication between clients and the SAP HANA system.</p> <p><code>global</code> - Binds the processes to all interfaces. This option does not require an internal network address entry.</p> <p><code>internal</code> - Binds the processes to this address only and to all local host interfaces. This option requires an internal network address entry.</p> <p><code>local</code> - Opens the communication ports for internal usage on the local interfaces. This configuration is only an option for single-host installations as the server is not reachable from outside. This option does not require an internal network address entry.</p> <p>If you define a value other than <code>local</code>, the local interfaces will always be open.</p>
<i>Internal Network Address</i>	Specifies the internal subset address in CIDR notation.

- Review the summary, and select *Run* to finalize the configuration.

You can find more information about SAP HANA system internal network and the network security recommendations, in the *SAP HANA Master*, *SAP HANA Security Guide*, and the Network Administration section of this *SAP HANA Administration Guide*.

Results

You have configured the inter-service communication of an SAP HANA system. The parameter values are entered in the `global.ini` configuration file under `[communication]`.

Related Information

[Network Administration \[page 1040\]](#)

10.3.4.1.2 Configure SAP HANA Inter-Service Communication Using the Command-Line Interface

To prevent unauthorized access to the SAP HANA system via the internal communication channels in multiple-host systems, you can configure inter-service communication using the SAP HANA database lifecycle manager command-line interface.

Prerequisites

- The SAP HANA system has been installed or updated with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on with the required root user or system administrator user `<sid>adm` credentials.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblc
```

By default, `<sapmnt>` is `/hana/shared`.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblc
```

3. Select the index for the `configure_internal_network` action. Then select `Enter`.
4. Define the required parameters.

Field Name	Description
<i>Inter-Service Communication</i>	<p>Specifies the listen interface for the internal network communication between the services of an SAP HANA system. This is not related to the communication between clients and the SAP HANA system.</p> <p><code>global</code> - Binds the processes to all interfaces. This option does not require an internal network address entry.</p> <p><code>internal</code> - Binds the processes to this address only and to all local host interfaces. This option requires an internal network address entry.</p> <p><code>local</code> - Opens the communication ports for internal usage on the local interfaces. This configuration is only an option for single-host installations as the server is not reachable from outside. This option does not require an internal network address entry.</p> <p>If you define a value other than <code>local</code>, the local interfaces will always be open.</p>

Field Name	Description
<i>Internal Network Address</i>	Specifies the internal subset address in CIDR notation.

For more information about parameters for the `configure_internal_network` action, see [Related Information](#).

- Review the summary, and select `y` to finalize the configuration.

Results

You have configured the inter-service communication of an SAP HANA system. The parameter values are entered in the `global.ini` configuration file under `[communication]`.

This configuration task can also be performed in batch mode and using a configuration file. For more information about the available configuration methods, see *Using the SAP HANA Platform LCM Tools*.

Note

When using the command line, the options can be set interactively during configuration only if they are marked as interactive in the help description. All other options have to be specified in the command line. To call the help, in the SAP HANA resident HDBLCM directory of the SAP HANA system, execute the following command:

```
./hdblcsm --action=configure_internal_network --help
```

You can find more information about SAP HANA system internal network and the network security recommendations, in *SAP HANA Master Guide*, *SAP HANA Security Guide*, and in the scaling SAP HANA information in the Network Administration section of this *SAP HANA Administration Guide*.

Example

The following example configures the internal network communication with internal interface:

```
./hdblcsm --action=configure_internal_network --listen_interface=internal --
internal_address=10.66.8/21
```

Related Information

[Using the SAP HANA Platform LCM Tools \[page 921\]](#)

[Configuring the Network for Multiple Hosts \[page 1438\]](#)

[Network Administration \[page 1040\]](#)

[nstart \[page 198\]](#)

10.3.4.1.3 Configure SAP HANA Inter-Service Communication Using the Web User Interface

To prevent unauthorized access to the SAP HANA system via the internal communication channels in multiple-host systems, you can configure inter-service communication using the SAP HANA database lifecycle manager Web user interface.

Prerequisites

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
 - Microsoft Windows
 - Internet Explorer - Version 9 or higher
If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing **Tools > Compatibility View Settings**.
 - Microsoft Edge
 - Mozilla Firefox - Latest version and Extended Support Release
 - Google Chrome - Latest version
 - SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
 - Mac OS - Safari 5.1 or higher

Note

For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser:

Option	Description
	https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html
	<p>i Note</p> <p>The URL is case sensitive. Make sure you enter upper and lower case letters correctly.</p>

SAP HANA cockpit	<ol style="list-style-type: none"> 1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. https://<host_FQDN>:<port> <p>i Note</p> <p>FQDN = fully qualified domain name</p> <ol style="list-style-type: none"> 2. Drill down on the name of the system from <i>My Resources</i> or from a group. 3. The links in <i>Platform Lifecycle Management</i> each launch additional functionality, giving you expanded capabilities for managing the resource.
-------------------------	---

2. Select the *Configure Inter-Service Communication* tile.
3. Optional: Modify the following parameters in the *Advanced Parameters Configuration* dialog. To access the *Advanced Parameters Configuration* dialog, click on the gear icon in the footer bar of the SAP HANA HDBLCM Web user interface.

Option	Description
nostart	Prevents the SAP HANA system from being started.
Timeouts	Sets customized timeouts (start_instance, start_service, stop_instance, stop_service).

4. Provide the password of the <sid>adm user, then select *Next*.
5. Specify values for the following fields:

Field Name	Description
<i>Inter-Service Communication</i>	<p>Specifies the listen interface for the internal network communication between the services of an SAP HANA system. This is not related to the communication between clients and the SAP HANA system.</p> <p><code>global</code> - Binds the processes to all interfaces. This option does not require an internal network address entry.</p> <p><code>internal</code> - Binds the processes to this address only and to all local host interfaces. This option requires an internal network address entry.</p> <p><code>local</code> - Opens the communication ports for internal usage on the local interfaces. This configuration is only an option for single-host installations as the server is not reachable from outside. This option does not require an internal network address entry.</p> <p>If you define a value other than <code>local</code>, the local interfaces will always be open.</p>
<i>Internal Network Address</i>	Specifies the internal subset address in CIDR notation.

6. Review the summary, and select *Run* to finalize the configuration.

You can find more information about SAP HANA system internal network and the network security recommendations, in the *SAP HANA Master*, *SAP HANA Security Guide*, and the Network Administration section of this *SAP HANA Administration Guide*.

Results

You have configured the inter-service communication of an SAP HANA system. The parameter values are entered in the `global.ini` configuration file under `[communication]`.

Related Information

[SAPUI5 Developer Guide](#)

[Add an SAP HANA System \[page 122\]](#)

[Network Administration \[page 1040\]](#)

10.3.4.1.4 Monitoring the Network Between Multiple Hosts

For scale-out systems, it is possible to monitor network traffic between hosts using the *Monitor Network* link in the SAP HANA cockpit.

On the *Network Overview* page, you can view the number of hosts and use the following tabs to monitor the network for multiple hosts:

- *Network Traffic*
Use this tab to understand the role of each host and the size of the sent (*Request Size*) and received data (*Response Size*) between the hosts of the scale-out SAP HANA database. The sender host sends requests to the receiver host which responds. You can change the unit on the top right.
- *Network Speed Check (Internal Communication)*
The list offers an overview of all network channels between the involved hosts starting with the slowest network connection.
The *Measure Network Speed* link offers the possibility to measure the network speed between the hosts in a scale-out SAP HANA database. You can select the size of the package for the speed check.
- *Network Speed Check (System Replication Communication)*
The list offers an overview of all network channels between the involved hosts in the system replication configuration.
The *Measure Network Speed* link offers the possibility to measure the network speed between the hosts in a system replication configuration.

10.3.5 Scaling SAP HANA Extended Application Services, Classic Model

If you have an application based on SAP HANA XS classic, you can configure multiple SAP HANA XS instances to work in a scale out SAP HANA system.

Context

If you are expecting a high degree of concurrency, you may want to distribute the XS classic server across the various hosts in your system. This is not enabled at the system level by default. You can manually change this setting to the system level in the Administration editor by performing the following steps.

Procedure

1. In the Administration editor, choose the *Configuration* tab.
2. Enter the string "instances" in the Filter box.
This search string returns a list of instances.
3. Change the instances setting to system level
 - a. Set the value of instances to "1" on the system level for `xsengine` and `sapwebdisp`.
4. Clear any entries on host level
 - a. Right-click on the green circle and choose *Delete* to clear any entries on the host level
 - b. In the *Delete Configuration Value* dialog box select the check box beside *HOST* layer and choose *Delete*.

Related Information

[Change a System Property in SAP HANA Studio \[page 301\]](#)

[Configure HTTP Load Balancing for SAP HANA Extended Application Services, Classic Model \[page 1211\]](#)

10.3.6 Starting and Stopping Distributed SAP HANA Systems Using SAPControl

You can use `SAPControl` to start or stop all the hosts in a scaled-out SAP HANA system from the command line.

i Note

You must be logged on to the SAP system host as user `<sid>adm` or as a user with root permissions.

Action	Command
Start the system	<code>/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StartSystem HDB</code>
Stop the system	<code>/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StopSystem HDB</code>
Query current status of all hosts in the system	<code>/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function GetSystemInstanceList</code>

i Note

HDB start or HDB stop only starts and stops the local host.

11 SAP HANA Deployment Infrastructure

An overview of how to set up and maintain the SAP HANA Deployment Infrastructure.

The SAP HANA Deployment Infrastructure provides a service that enables you to deploy database development artifacts to so-called containers. This service includes a family of consistent design-time artifacts for all key SAP HANA platform database features which describe the target (run-time) state of SAP HANA database artifacts, for example: tables, views, or procedures. These artifacts are modeled, staged (uploaded), built, and deployed into SAP HANA.

! Restriction

The HDI focuses strictly on deployment; HDI does not include any version-control tools, nor does it provide any tools for life-cycle management.

HDI provides its services using a separate database process named `diserver`. On systems where XS advanced is installed, HDI is already enabled; on other systems where XS advanced is not installed, the `diserver` process must usually be enabled by the database administrator before HDI can be used. If required by the usage scenario, other database process may also need to be started as well.

The SAP HANA service broker is used to create and drop HDI containers; each HDI container comprises a design-time container (DTC) and a run-time container (RTC). The HDI deployment tools deploy database artifacts to an HDI container. Design-time database objects are typically located in the `db/` folder of the application design-time hierarchy. The deployment process populates the database run-time with the specified catalog objects. In addition to database artifacts, HDI also enables you to import and export table content such as business configuration data and translatable texts.

! Restriction

HDI enables you to deploy database objects only; it is not possible (or necessary) to use HDI to deploy application-layer artifacts such as JavaScript programs or OData objects.

The configuration of HDI containers also involves the creation and configuration of the following design-time artifacts:

- Container deployment configuration (`.hdiconfig`)
A JSON file containing a list of the bindings between database artifact types (for example, sequence, procedure, table) and the corresponding plug in (and version) required to deploy the artifacts.
- Run-time container namespace rules (`.hdinamespace`)
An **optional** JSON file containing a list of design-time file suffixes and the naming rules for the corresponding run-time locations.

→ Tip

You can apply the rules defined in the `.hdinamespace` file either exclusively to the folder it is located in or to the folder it is located in and its subfolders.

HDI Maintenance Tasks

Maintaining the HDI, its containers, and container groups, involves the following high-level tasks:

- **Enabling the HDI**
A database administrator with SYSTEM privileges starts the HDI for the first time, creates the necessary administrator users, and assigns the new users the access privileges required to administrate the HDI.
- **Maintaining the HDI**
An HDI administrator configures HDI, creates HDI container groups, and grants and revokes the access privileges required by the HDI container-group administrators.
- **Maintaining HDI container groups**
HDI container-group administrators drop and create HDI containers, grant and revoke container (and container-group) access privileges, and import and export containers (for support purposes).
- **Maintaining HDI containers**
HDI container administrators grant and revoke container-based access privileges, configure libraries and parameters, and grant and revoke access to a HDI container's schemas.

Related Information

[HDI Administrator Roles \[page 1454\]](#)

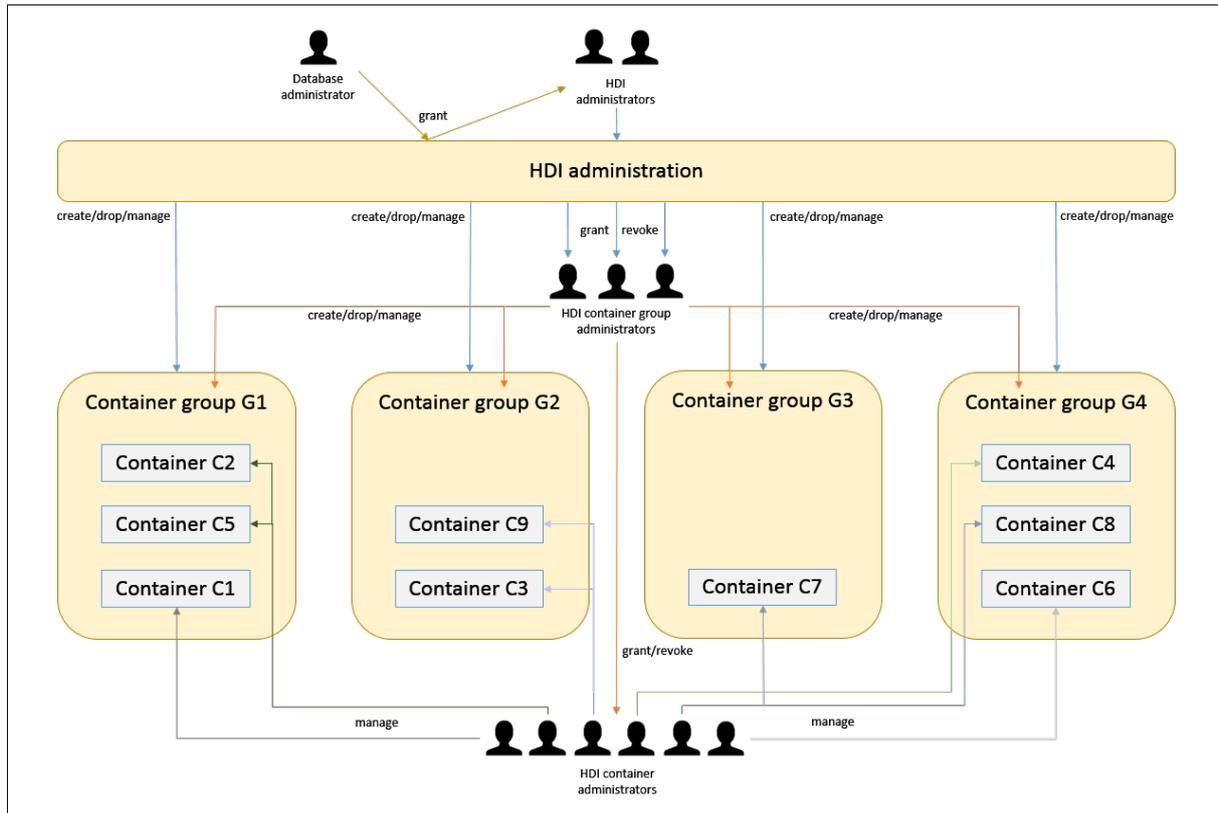
[HDI Administration in Context \[page 1452\]](#)

[List All Currently Configured Build Plug-in Libraries Available to a Container \[page 1519\]](#)

11.1 HDI Administration in Context

An overview of the HDI administration process including the administrator users who set up and maintain HDI and its components.

The relationships between the various roles involved in the maintenance of HDI containers and the corresponding scope of each role is shown in the following diagram:



HDI Administration

HDI provides its services using a separate database process named `diserver`. On systems where XS advanced is installed, HDI is already enabled; on other systems where XS advanced is not installed, the `diserver` process must usually be enabled by the database administrator before HDI can be used. If required by the usage scenario, other database process may also need to be started as well.

The database administrator `SYSTEM` is needed for the first-time enabling of HDI in SAP HANA and for creating an HDI administrator, who then performs the tasks required to set up and maintain other HDI administrators, if required.

i Note

`SYSTEM` user privileges are required to create the first HDI administrator, who can then create other HDI administrators. After creation of the first HDI administrator, the `SYSTEM` user can be deactivated.

Created by the database administrator, the HDI administrator is responsible for the setup and overall maintenance of HDI. The role of the HDI administrator includes configuring general HDI parameters,

maintaining containers and container groups (for example, by creating and dropping containers and container groups), and managing container-group administrator privileges, for example, by granting and revoking HDI container-group access permissions.

The HDI container-group administrator manages a set of containers in container groups assigned by the HDI administrator. Container-group management tasks include: granting and revoking container (and container-group) **administrator** access privileges; granting and revoking container **user** access privileges (for example, for temporary support purposes); maintaining containers and container groups.

→ Tip

The APIs of a container group “G” are in the `_SYS_DI#G` schema.

The HDI container administrator manages one or more containers assigned by the container-group administrator. The role of the container-manager focuses primarily on configuring and controlling access to the HDI containers used to store the database objects deployed by the SAP HANA Deployment Infrastructure deploy service and repairing any problems that occur with run-time objects in the assigned HDI containers. An HDI container administrator can manage one or more containers in one HDI container group or multiple containers distributed across multiple container groups.

→ Tip

The APIs of a container “C” are in the `C#DI` schema.

HDI Containers

The SAP HANA Deployment Infrastructure (HDI) provides a service that enables you to deploy database development artifacts to so-called containers. This service includes a family of consistent design-time artifacts for all key HANA platform database features which describe the target (run-time) state of SAP HANA database artifacts, for example: tables, views, or procedures. These artifacts are modeled, staged (uploaded), built, and deployed into SAP HANA.

i Note

The HDI focuses strictly on deployment; HDI does not include any version-control tools, nor does it provide any tools for life-cycle management.

The SAP HANA service broker is used to create and destroy HDI containers; each HDI container comprises a design-time container (DTC), which is an isolated environment used to store design-time files, and a run-time container (RTC), which is used to store deployed objects built according to the specification stored in the corresponding design-time artifacts.

The deployment process populates the database run-time with the specified catalog objects. In addition to database artifacts, HDI also enables you to import and export table content such as business configuration data and translatable texts.

! Restriction

HDI enables you to deploy database objects only; it is not possible (or necessary) to deploy application-layer artifacts such as JavaScript programs or OData objects.

The configuration of HDI containers also involves the creation and configuration of the following design-time artifacts:

- Container deployment configuration (`.hdi.config`)
A JSON file containing a list of the bindings between database artifact types (for example, sequence, procedure, table) and the corresponding deployment plug in (and version).
- Run-time container namespace rules (`.hdi.namespace`)
A JSON file containing a list of design-time file suffixes and the naming rules for the corresponding runtime locations.

→ Tip

You can apply the rules defined in the `.hdi.namespace` file either exclusively to the folder it is located in or to the folder it is located in **and** its subfolders.

HDI Container Groups

HDI container groups are logical collections of the HDI containers used to store the objects deployed by the SAP HANA Deployment Infrastructure deploy service. After creation, an HDI container group can be assigned to a dedicated container-group administrator, who must be granted the privileges required to perform the typical tasks associated with the administration of container group, for example: granting and revoking container (and container-group) access privileges; and maintaining container groups (and the containers assigned to the groups).

→ Tip

Container groups are intended to make life easier when multiple administrators require access to containers from different contexts, for example, XS advanced, ABAP, or other development groups working in native SAP HANA contexts.

The HDI administrator can create container groups for a target audience whose container-related requirements are unique or where there is an obvious benefit for a logical separation, for example, between ABAP and applications running in the XS advanced model run-time environment.

Related Information

[Enabling HDI in the Database \[page 1460\]](#)

[Maintaining the HDI \[page 1465\]](#)

[Maintaining HDI Container Groups \[page 1494\]](#)

[Maintaining HDI Containers \[page 1508\]](#)

11.2 HDI Administrator Roles

HDI administration involves a number of tasks that must be performed by different administrator roles.

The following table describes the scope of the various HDI administrator roles and lists the most common tasks the administrators are expected to perform.

HDI Admin Scope, Roles, and Tasks

HDI Role	Description	Common Tasks
Database administrator	<p>The database administrator SYSTEM is needed for initially enabling HDI and for creating an HDI administrator.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>SYSTEM user privileges are required to create the first HDI administrator, who can then create other HDI administrators. After creation of the first HDI administrator, the SYSTEM user can be deactivated.</p> </div>	<ul style="list-style-type: none"> Enable HDI Create an HDI administrator Grant and revoke HDI administrator privileges
HDI administrator	<p>Configures general HDI parameters, maintains containers and container groups, and manages container group administrator privileges.</p>	<ul style="list-style-type: none"> Configure HDI Create and drop containers and container groups Grant and revoke required access privileges Maintain containers and container groups Move containers between container groups
HDI container-group administrator	<p>Manages the container groups assigned by the HDI administrator. The APIs of a container group "G" are in the <code>_SYS_DI#G</code> schema.</p>	<ul style="list-style-type: none"> Grant and revoke container (and container-group) administrator access privileges Import and export containers (for support purposes) Grant and revoke container user access privileges (for support purposes) Maintain containers and container groups
HDI container administrator	<p>Configures and controls access to a container and manages run-time objects in the assigned containers. The APIs of a container "C" are in the <code>C#DI</code> schema.</p>	<ul style="list-style-type: none"> Grant and revoke container administrator access privileges Configure libraries and parameters Grant and revoke roles from schemas to users Grant and revoke user access to container schemas Cancel an asynchronous make operation

Related Information

[SAP HANA Deployment Infrastructure \(Administration Guide\) \[page 1450\]](#)

11.3 The SQL API for SAP HANA Deployment Infrastructure (HDI)

An SQL application programming interface (API) is available to help maintain the SAP HANA Deployment Infrastructure (HDI).

In this topic, you can find information about the following components of the SQL API for HDI:

- [SQL APIs for HDI \[page 1456\]](#)
- [Parameters for API Procedures \[page 1457\]](#)
- [Table Types \[page 1457\]](#)
- [Predefined Parameter Tables \[page 1458\]](#)
- [Result Sets \[page 1458\]](#)
- [Return Codes \[page 1458\]](#)
- [Request ID \[page 1458\]](#)
- [Messages \[page 1459\]](#)

SQL APIs for HDI

SAP HANA Deployment Infrastructure can be seen as a layer on top of the SAP HANA database. The HDI offers an SQL-based API that is accessible by means of standard SAP HANA SQL connection data. The following APIs are provided to help you manage and work with HDI:

SQL APIs for HDI

HDI API	Description
HDI administration API	Used mainly for managing container groups and the administrative access to them. Use of the HDI administration API requires the privileges of an HDI administrator. i Note The privileges of the SYSTEM user are required to create the first HDI administrator, who can then create other HDI administrators, if required. After creation of the first HDI administrator, the SYSTEM user can be deactivated.
HDI container group administration API	Used for managing a set of containers inside a container group and the administrative access to them. To use this API, a user needs to be granted the privileges of an HDI container group administrator by an HDI administrator.
HDI container administration API	Used for configuring a container and controlling the access to it. To use this API, a user needs to be granted the privileges of an HDI container administrator by an HDI container group administrator.

HDI API	Description
HDI container content development API:	Used for enabling applications to deploy or undeploy HDI artifacts within a container and manage access to the containers. To use this API, a user must be granted the privileges of an HDI container content developer by an HDI container administrator.

i Note

For more information about the HDI container content development API, see the *SQL API for SAP HANA Deployment Infrastructure (HDI)* in *Related Information* below.

→ Tip

For more information about the HDI Container APIs, see *Related Information*.

Parameters of HDI API Procedures

HDI APIs usually expect a number of input and output parameters, where output parameters are always being the last parameters to be passed. Most of the HDI APIs return the following information as output parameters:

- A return code
- A request ID
- A messages table

Table Types

HDI's predefined table types can be found in the `_SYS_DI` schema; they can be identified by the prefix "TT_" "grant" APIs available to the HDI administrator, the HDI container-group administrator, or the HDI container administrator. To pass a parameter to an HDI API that expects an HDI table type, a temporary table of that type must be created before calling the API procedure. This temporary table must then be filled with the data that the API expects. For example:

≡ Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #CONFIG_PARAMETERS LIKE
_SYS_DI.TT_PARAMETERS;
INSERT INTO #CONFIG_PARAMETERS (KEY, VALUE) VALUES ('make.max_parallel_jobs',
'8');
CALL C#DI.CONFIGURE_CONTAINER_PARAMETERS(#CONFIG_PARAMETERS,
_SYS_DI.T NO_PARAMETERS, ?, ?, ?);
DROP TABLE #CONFIG_PARAMETERS;
```

Predefined Parameter Tables

HDI's predefined tables can be found in the `_SYS_DI` schema; they can be identified by the prefix "T_". The privileges required to access these tables are granted to a user by means of the grant APIs available to the HDI administrator, the HDI container-group administrator, or the HDI container administrator. The predefined tables provide tables with default content that are needed more frequently. For example, when no special parameters need to be given to an HDI API, the `_SYS_DI.T_NO_PARAMETERS` table can be used.

Sample Code

```
CALL C#DI.CONFIGURE_CONTAINER_PARAMETERS(_SYS_DI.T_NO_PARAMETERS,  
_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

Result Sets

The SAP HANA DI SQL API calls usually return the following result sets;

1. A result set that contains a result code (for example, 0 or 1) and a request ID (a unique ID for the API call).
2. A result set that includes a table of messages containing information about the execution of the HDI call.

Return Codes

The return code indicates if the API call was successful, or if there were problems during the execution of the call. The meaning of the possible return codes is described in the following table:

SQL API Return Codes for HDI

Return Code	Meaning	Details
0	Success	The call was successful, and the messages did not report any warnings or errors
1	Warning	The messages contain warnings, but no errors.
-1	Error	The messages contain errors.
-2	Fatal Error	No messages could be logged. This indicates a problem with the database.

Request ID

The request ID is a unique ID generated for each API call. This ID is always the same for messages that originate from the same API call. It can, for example, be used by a container administrator to query the messages that were produced by a certain API call.

Messages

Most APIs return a table of messages with information about the execution of the HDI call. The message, content, and format is explained in the following table:

SQL API Message Tables for HDI

Message Table	Description	Optional?
REQUEST_ID	The unique ID of the API call that produced this message. This ID is always the same for messages that originate from the same API call.	
ROW_ID	An increasing number representing the line number of the message log.	
LEVEL	The indentation level of the message (used for better visual representation).	
TYPE	The type of message returned; the following list describes the possible message types: <ul style="list-style-type: none">• SUMMARY A summary of the API call• HDI: The message from HDI• PLUGIN: The message from a plug-in	
LIBRARY_ID	The ID of the library (for messages from a plug-in)	Yes
PLUGIN_ID	The ID of the plug-in (for messages from a plug-in)	Yes
PATH	The path to the artifact that is being processed	Yes
SEVERITY	The severity of the message, for example: INFO, WARNING, ERROR, ...	
MESSAGE_CODE	A unique code corresponding to the MESSAGE field	
MESSAGE	The message text	
LOCATION	The Position (for example, "line:column") within the artifact that the message refers to	Yes
LOCATION_PATH	The XPath expression within the artifact that the message refers to	Yes
TIMESTAMP_UTC	The time stamp indicating when the message was created	

Related Information

[SAP HANA Deployment Infrastructure \[page 1450\]](#)

[The SQL API for SAP HANA Deployment Infrastructure \(HDI\) \[page 1456\]](#)

11.4 Enabling HDI in the Database

Start the SAP HANA Deployment Infrastructure (HDI) services for the first time.

HDI provides its services using a separate database process, `diserver`. If SAP HANA extended application services, advanced model (XS advanced), is installed in the system, then HDI is already enabled. Otherwise, a database administrator must enable HDI manually, for example, by starting the `diserver` process, before HDI can be used.

i Note

Depending on your scenario, further database processes may also need to be started.

Enabling HDI typically involves the following administrator tasks:

- Enable and disable HDI in the database *
- Create an HDI administrator *
- Revoke HDI administrator privileges

i Note

The SYSTEM user is required to enable HDI and create an initial HDI administrator, who can then configure HDI and create additional HDI administrators if required. If the SYSTEM user is deactivated (recommended), you will need to reactivate it temporarily to create the first HDI administrator, for example, with the SQL statement: `ALTER USER SYSTEM ACTIVATE USER NOW`. After you have completed these setup tasks, you can deactivate the SYSTEM user again and use the HDI administrator user to perform setup tasks.

Related Information

[Enable HDI for a Specific Tenant on a Multi-Tenant Database \[page 1461\]](#)

[Create an HDI Administrator \[page 1463\]](#)

[Revoke the HDI Administrator Privileges \[page 1464\]](#)

11.4.1 Enable HDI for a Specific Tenant on a Multi-Tenant Database

If SAP HANA XS advanced model (XS advanced) is not installed in your system, you must enable HDI in the relevant tenant database.

Prerequisites

The SYSTEM user has been reactivated and you have its credentials.

i Note

SYSTEM user privileges are required to enable HDI. After enabling HDI, the SYSTEM user can be deactivated.

Context

HDI provides its services using a separate database process called `diserver`. On systems where XS advanced is installed, HDI is already enabled; on other systems where XS advanced is not installed, the `diserver` process must usually be enabled by the database administrator before HDI can be used. Additionally, if required by the usage scenario, other database process may need not be started as well.

To enable HDI for a specific tenant on a multi-tenant database, perform the following steps:

Procedure

1. In an SQL console, connect to the `systemDB` database as the SYSTEM user.
2. Enable HDI in the target tenant database.

Insert the following SQL code into the SQL console:

→ Tip

For the `dbName` declaration in the following example, replace `XY1` with the name of the tenant database for which HDI should be enabled.

```
DO
BEGIN
  DECLARE dbName NVARCHAR(25) = 'XY1';
  DECLARE diserverCount INT = 0;
  DECLARE scriptserverCount INT = 0;
  DECLARE dpserverCount INT = 0;
  DECLARE docstoreCount INT = 0;
  -- Start diserver
```

```

SELECT COUNT(*) INTO diserverCount FROM SYS_DATABASES.M_SERVICES WHERE
SERVICE_NAME = 'diserver' AND DATABASE_NAME = :dbName AND ACTIVE_STATUS =
'YES';
IF diserverCount = 0 THEN
EXEC 'ALTER DATABASE ' || :dbName || ' ADD ''diserver''';
END IF;
-- [OPTIONAL] For AFLLang Procedure artifacts
SELECT COUNT(*) INTO scriptserverCount FROM SYS_DATABASES.M_SERVICES WHERE
SERVICE_NAME = 'scriptserver' AND DATABASE_NAME = :dbName AND ACTIVE_STATUS =
'YES';
IF scriptserverCount = 0 THEN
EXEC 'ALTER DATABASE ' || :dbName || ' ADD ''scriptserver''';
END IF;
-- [OPTIONAL] For Flow Graphs or Replication Task artifacts
SELECT COUNT(*) INTO dpserverCount FROM SYS_DATABASES.M_SERVICES WHERE
SERVICE_NAME = 'dpserver' AND DATABASE_NAME = :dbName AND ACTIVE_STATUS =
'YES';
IF dpserverCount = 0 THEN
EXEC 'ALTER DATABASE ' || :dbName || ' ADD ''dpserver''';
END IF;
-- [OPTIONAL] For JSON DocStore and hdbcollection artifacts
SELECT COUNT(*) INTO docstoreCount FROM SYS_DATABASES.M_SERVICES WHERE
SERVICE_NAME = 'docstore' AND DATABASE_NAME = :dbName AND ACTIVE_STATUS =
'YES';
IF docstoreCount = 0 THEN
EXEC 'ALTER DATABASE ' || :dbName || ' ADD ''docstore''';
END IF;
END;

```

- a. Configure support for (Application Function Library language) procedures, if required.

If support for “AFLLang” (Application Function Library language) procedures is required, use the optional section to start the `scriptserver` process, otherwise remove the section.

- b. Configure support for flow graphs or replication tasks, if required.

If support for flow graphs (`hdbflowgraph`) or replication tasks (`hdbreptask`) is required, use the optional paragraph to start the `dpserver` process, otherwise remove the section.

- c. Configure support for the JSON Document Store (DocStore), if required.

If support for the JSON Document Store (DocStore) and collection artifacts (`.hdbcollection`) is required, use the optional section to start the `docstore` process, otherwise remove the section.

3. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

4. Confirm that the desired processes are running.

i Note

Check which processes are running in the database administration view.

Related Information

[HDI Administrator Roles \[page 1454\]](#)

11.4.2 Create an HDI Administrator

Create an HDI administrator user. The HDI administrator is responsible for configuring general HDI parameters, creating and dropping container groups, moving containers between groups, and managing container group administrator privileges.

Prerequisites

- The SYSTEM user has been reactivated and you have its credentials.
- HDI is enabled.

Context

This method uses the predefined `_SYS_DI.T_DEFAULT_DI_ADMIN_PRIVILEGES` table, which contains the largest possible set of privileges that can be granted for a user of this type. It is also possible to reduce the set of privileges by explicitly specifying the desired set of privileges and not using this default table.

i Note

A variant of this procedure also exists, namely:

`_SYS_DI.GRANT_CONTAINER_GROUP_API_PRIVILEGES_WITH_GRANT_OPTION`; it grants the given privileges `WITH GRANT OPTION` to the target user. However, this variant procedure is only needed in special scenarios, for example, when building a custom SQL API by wrapping the HDI SQL API in SQLScript procedures.

Procedure

1. In an SQL console, connect to the database with the SYSTEM user.
2. Optional: Create a new user by executing the following statement:

```
CREATE USER <HDI_admin_username> PASSWORD <password> NO
FORCE_FIRST_PASSWORD_CHANGE;
```

3. Grant the new HDI administrator user the required privileges by executing the following statement:

```
CREATE LOCAL TEMPORARY TABLE #PRIVILEGES LIKE _SYS_DI.TT_API_PRIVILEGES;
INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME) SELECT
'<HDI_admin_username>', PRIVILEGE_NAME, OBJECT_NAME FROM
_SYS_DI.T_DEFAULT_DI_ADMIN_PRIVILEGES;
CALL _SYS_DI.GRANT_CONTAINER_GROUP_API_PRIVILEGES('_SYS_DI', #PRIVILEGES,
_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;
```

4. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

5. Confirm that the new HDI administrator user can call HDI API procedures in the `_SYS_DI` schema.

Next Steps

Deactivate the SYSTEM user: `ALTER USER SYSTEM DEACTIVATE USER NOW.`

Related Information

[Enable HDI for a Specific Tenant on a Multi-Tenant Database \[page 1461\]](#)

11.4.3 Revoke the HDI Administrator Privileges

Revoke the HDI administrator privileges from a specified user.

Prerequisites

- The SYSTEM user has been reactivated and you have its credentials.
- HDI is enabled.

Context

The database administrator SYSTEM can revoke the HDI administrator privileges from a user.

i Note

This method uses the predefined `_SYS_DI.T_DEFAULT_DI_ADMIN_PRIVILEGES` table, which contains the largest possible set of privileges that can be revoked from a user of this type. It is also possible to reduce the set of privileges by explicitly specifying the desired set of privileges and not using this default table.

Procedure

1. In an SQL console, connect to the database with the SYSTEM user.
2. Open the SQL editor for this database.

3. Revoke the HDI administrator privileges from the selected user.

Insert the following SQL code into the SQL console:

Note

In the following code example, replace `NEW_HDI_ADMIN` with the name of the user from whom the HDI administrator privileges should be revoked.

```
CREATE LOCAL TEMPORARY TABLE #PRIVILEGES LIKE _SYS_DI.TT_API_PRIVILEGES;
INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME) SELECT
'NEW_HDI_ADMIN', PRIVILEGE_NAME, OBJECT_NAME FROM
_SYS_DI.T_DEFAULT_DI_ADMIN_PRIVILEGES WHERE NOT (PRIVILEGE_NAME = 'SELECT'
AND OBJECT_NAME LIKE 'T%');
CALL _SYS_DI.REVOKE_CONTAINER_GROUP_API_PRIVILEGES('_SYS_DI', #PRIVILEGES,
_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;
```

4. Execute the SQL code.
Confirm that the SQL code completes successfully and displays the HDI return code 0.
5. (Optional) Confirm that the `NEW_HDI_ADMIN` user is no longer able to call HDI API procedures in the `_SYS_DI` schema.

Next Steps

Deactivate the SYSTEM user, for example, with the following command in the SQL console:

```
ALTER USER SYSTEM DEACTIVATE USER NOW
```

Related Information

[Enabling HDI in the Database \[page 1460\]](#)

11.5 Maintaining the HDI

Maintenance of the SAP HANA Deployment infrastructure is the responsibility of the HDI administrator, who must set up and configure general HDI parameters.

Managing the HDI typically involves the following administrator tasks:

- Configure HDI
- List plug-in libraries available to a container
- Create and drop containers and container groups
- Grant and revoke container-group administration privileges
- Maintain containers and container groups

- Move containers between container groups

Exporting and Importing Containers

In special cases, a container and its dependencies can be exported from the database and imported into a different database for support purposes.

⚠ Caution

The API procedures `_SYS_DI.EXPORT_CONTAINER_FOR_SUPPORT` and `_SYS_DI.IMPORT_CONTAINER_FOR_SUPPORT` are available to the HDI administrator but intended for use by SAP support, exclusively. Bear in mind that the exported data might include private or confidential data from the container, and the container import could also compromise the integrity of the database.

Related Information

[Configure HDI Parameters \[page 1469\]](#)

[Enabling HDI in the Database \[page 1460\]](#)

[Maintaining HDI Container Groups \[page 1494\]](#)

[Maintaining HDI Containers \[page 1508\]](#)

11.5.1 HDI Container Group Administration

HDI container-group administration involves the management of a set of containers in an assigned container group.

Most container-group-related administration tasks can be performed by both the HDI container-group administrator and the HDI administrator; the tasks include creating and dropping containers in the container group or, if necessary, granting container-administration privileges to other users of a container. However, only the HDI administrator can create and drop a container **group**. After creating a container group, the HDI administrator creates a new container-group administrator by assigning the necessary container-group administrator privileges to one of more selected users. The new container-group administrator is responsible for maintaining the containers in the new container group.

Each container group can have its own set of administrators. Administrative privileges for a container group can only be granted by an HDI administrator or a container group administrator who has the necessary privileges. Container-group administrators can only grant privileges for the container groups for which they are directly responsible.

For the HDI administrator, the APIs of a container group are located in the schema `_SYS_DI`; for the HDI container-group administrator, the APIs of a container group called "G" are located in the schema `_SYS_DI#G`. For more information about the functionality available to the HDI container-group administrator, see the section *Maintaining HDI Container Groups* listed in *Related Information*.

⚠ Caution

The administration of a container group should normally be performed by the container group's assigned administrator. The HDI administrator should only be used for container-group administration purposes in exceptional circumstances.

Exporting and Importing Containers

In special cases, a container and its dependencies can be exported from the database and imported into a different database for support purposes.

⚠ Caution

The API procedures `_SYS_DI#G.EXPORT_CONTAINER_FOR_SUPPORT` and `_SYS_DI#G.IMPORT_CONTAINER_FOR_SUPPORT` are available to the HDI container-group administrator but intended for use only by SAP support. The exported data might include private or confidential data from the container, and the container import could also compromise the integrity of the database.

The same API procedures are also available to the HDI administrator, but in the schema `_SYS_DI`, for example, `_SYS_DI.EXPORT_CONTAINER_FOR_SUPPORT`.

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

11.5.2 HDI Container Administration

Most container-related administration tasks can be performed by both the HDI administrator and the HDI container administrator.

Container-related administrator tasks are performed by calling the respective HDI container-administration SQL procedures, not of a target container, but of the `_SYS_DI` schema of the HDI administrator with an additional first parameter taking the name of the target container. For details about the tasks required to maintain HDI containers and the functionality available to help perform these tasks, refer to the section about *Maintaining HDI Containers* listed in *Related Information*.

To grant a user "U" the privileges to access the run-time objects in the container "C", the HDI administrator can call the SQL procedure `GRANT_CONTAINER_SCHEMA_PRIVILEGES` in the `_SYS_DI` schema, passing the container's name (in this example, "C") as the additional first parameter, as illustrated in the following example:

⚠ Caution

The administration of a container should normally be performed by the container's assigned administrator. The HDI administrator should only be used for this purpose in exceptional circumstances.

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
_SYS_DI.TT_SCHEMA_PRIVILEGES;
INSERT INTO #PRIVILEGES ( PRIVILEGE_NAME, PRINCIPAL_SCHEMA_NAME,
PRINCIPAL_NAME ) VALUES ( 'SELECT', '', 'U' );
CALL _SYS_DI.GRANT_CONTAINER_SCHEMA_PRIVILEGES( 'C', #PRIVILEGES,
_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;
```

Related Information

[Maintaining HDI Containers \[page 1508\]](#)

11.5.3 HDI Container Schemas

An HDI container makes use of multiple schemas; the different schemas serve different aims and tasks.

Maintaining HDI containers involves the configuration and use of the schemas listed and described in the following table:

Note

In the following table, the schema names are based on the assumption that the base HDI container is named "C".

HDI Container Schema Names and Usage

Container Schema Name	Description
C	Contains generated database objects that belong to a special object-owner user called c#00. The database objects are generated from design-time objects in container C.
C#DI	Contains the API procedures and HDI-internal data required for the container management
C#00	The schema for the user to whom the artifacts in the base container "C" belong. The user schema C#00 is empty.

Related Information

[Maintaining HDI Containers \[page 1508\]](#)

[Configure HDI Parameters \[page 1469\]](#)

[Maintaining the HDI \[page 1465\]](#)

11.5.4 Configure HDI Parameters

The HDI administrator can configure some general aspects of HDI with configuration parameters, for example, how long an HDI operation waits for a locking conflict to clear or default behavior of containers.

Procedure

1. In an SQL console, connect to the database as the HDI administrator.
2. Configure the required configuration parameters by executing the following SQL statement:

```
CREATE LOCAL TEMPORARY COLUMN TABLE #CONFIG_PARAMETERS LIKE
SYS_DI.TT_PARAMETERS;
INSERT INTO #CONFIG_PARAMETERS (KEY, VALUE) VALUES
('make.default_max_parallel_jobs', '<value>');
-- insert more parameters as required
CALL SYS_DI.CONFIGURE_DI_PARAMETERS(#CONFIG_PARAMETERS,
SYS_DI.TT_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #CONFIG_PARAMETERS;
```

For a description of all available parameters, see *SAP HANA DI Configuration Parameters*.

3. Execute the SQL code.
Confirm that the SQL code completes successfully and displays the HDI return code 0.
4. Verify that the configuration parameters have been set correctly in the `diserver.ini` configuration file.

Related Information

[SAP HANA DI Configuration Parameters \[page 1484\]](#)

11.5.4.1 SAP HANA DI Parameters

Overview of available SAP HANA DI and build-plugin parameters.

In SAP HANA Deployment Infrastructure (HDI), parameters are a means of controlling the execution flow of SAP HANA DI procedure calls. SAP HANA DI includes the following parameter types:

- SAP HANA DI parameters
SAP HANA DI parameters are used to control the execution flow of SAP HANA DI procedures and SAP HANA DI container-specific procedures. For example, they specify the time a container operation waits for a locking conflict to clear or they indicate if warnings during an SAP HANA DI call should be treated as errors.
- Build-plugin parameters
Build plug-in parameters control the execution flow of the deployment process of a build plug-in for all database objects of the corresponding type. For example, a build-plug-in parameter can be used to specify the batch size for batched database access or for batch-processing within a build plug-in.

SAP HANA DI Procedures

The following table lists the available parameters for SAP HANA DI procedures.

SAP HANA DI Call	Available Parameters
<code>_SYS_DI.CANCEL</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> treat_warnings_as_errors message_severity
<code>_SYS_DI.CONFIGURE_CONTAINER_PARAMETERS</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> message_severity
<code>_SYS_DI.CONFIGURE_DI_PARAMETERS</code>	trace_context trace_level.<trace topic> message_severity
<code>_SYS_DI.CONFIGURE_LIBRARIES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> undeploy message_severity
<code>_SYS_DI.CREATE_CONTAINER</code>	trace_context trace_level.<trace topic> message_severity
<code>_SYS_DI.CREATE_CONTAINER_GROUP</code>	trace_context trace_level.<trace topic> message_severity
<code>_SYS_DI.DROP_CONTAINER</code>	container_lock_wait_timeout ignore_deployed ignore_errors ignore_work trace_context trace_level.<trace topic> message_severity

SAP HANA DI Call

Available Parameters

`_SYS_DI.DROP_CONTAINER_GROUP`

move_containers_to_default_group
group
trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.EXPORT_CONTAINER *`

⚠ Caution

Removed with SAP HANA 2.0 SPS 00. See `_SYS_DI.EXPORT_CONTAINER_FOR_SUPPORT`

container_lock_wait_timeout
trace_context
trace_level.<trace topic>

`_SYS_DI.EXPORT_CONTAINER_FOR_SUPPORT`

container_lock_wait_timeout
export_container_schema_data
export_container_schema_foreign_objects
trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.GRANT_CONTAINER_API_PRIVILEGES`

container_lock_wait_timeout
trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.GRANT_CONTAINER_API_PRIVILEGES_WITH_GRANT_OPTION`

container_lock_wait_timeout
trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.GRANT_CONTAINER_GROUP_API_PRIVILEGES`

trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.GRANT_CONTAINER_GROUP_API_PRIVILEGES_WITH_GRANT_OPTION`

trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.GRANT_CONTAINER_SCHEMA_PRIVILEGES`

container_lock_wait_timeout
trace_context
trace_level.<trace topic>
message_severity

SAP HANA DI Call

Available Parameters

`_SYS_DI.GRANT_CONTAINER_SCHEMA_ROLES`

container_lock_wait_timeout
trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.GRANT_CONTAINER_SUPPORT_PRIVILEGE`

container_lock_wait_timeout
trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.IMPORT_CONTAINER *`

⚠ Caution

Removed with SAP HANA 2.0 SPS 00. See `_SYS_DI.IMPORT_CONTAINER_FOR_SUPPORT`.

container_lock_wait_timeout
trace_context
trace_level.<trace topic>

`_SYS_DI.IMPORT_CONTAINER_FOR_SUPPORT`

container_lock_wait_timeout
accept_risk_of_database_corruption_by_container_import
trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.LIST_CONFIGURED_LIBRARIES`

container_lock_wait_timeout
trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.LIST_LIBRARIES`

trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.MOVE_CONTAINER_TO_GROUP`

trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.REVOKE_CONTAINER_API_PRIVILEGES`

container_lock_wait_timeout
trace_context
trace_level.<trace topic>
message_severity

SAP HANA DI Call

Available Parameters

`_SYS_DI.REVOKE_CONTAINER_GROUP_API_PRIVILEGES`

trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.REVOKE_CONTAINER_SCHEMA_PRIVILEGES`

container_lock_wait_timeout
trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.REVOKE_CONTAINER_SCHEMA_ROLES`

container_lock_wait_timeout
trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.REVOKE_CONTAINER_SUPPORT_PRIVILEGE`

container_lock_wait_timeout
trace_context
trace_level.<trace topic>
message_severity

`_SYS_DI.CONFIGURE_CONTAINER *`

⚠ Caution

Deprecated since SAP HANA 1.0 SPS 12.

`_SYS_DI.CONFIGURE_DI *`

⚠ Caution

Deprecated since SAP HANA 1.0 SPS 12.

Example: Calling an SAP HANA DI Procedure with Parameters Set

📄 Sample Code

```
-- prepare parameters table
create table MY_PARAMETERS like _SYS_DI.TT_PARAMETERS;
insert into MY_PARAMETERS (KEY, VALUE) values ('ignore_work', 'true');
insert into MY_PARAMETERS (KEY, VALUE) values ('ignore_deployed', 'true');
-- call procedure
call _SYS_DI.DROP_CONTAINER('MY_CONTAINER', MY_PARAMETERS, ?, ?, ?);
```

SAP HANA DI Container-Specific Procedures

The following table lists the available parameters for SAP HANA DI container-specific procedures.

SAP HANA DI Container-Specific Call**Available Parameters**

<code><container>#DI.CANCEL</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> treat_warnings_as_errors message_severity
<code><container>#DI.CONFIGURE_CONTAINER_PARAMETERS</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> message_severity
<code><container>#DI.CONFIGURE_LIBRARIES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> undeploy message_severity
<code><container>#DI.DELETE</code>	container_lock_wait_timeout ignore_non_existing_paths recursive trace_context trace_level.<trace topic> message_severity
<code><container>#DI.GET_DEPENDENCIES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> variant message_severity
<code><container>#DI.GET_MAKE_GROUPS</code>	container_lock_wait_timeout max_parallel_jobs optimized_redeploy simulate_make skip_unchanged_expansions trace_context trace_level.<trace topic> treat_warnings_as_errors undeploy_dependent_recursively message_severity
<code><container>#DI.GRANT_CONTAINER_API_PRIVILEGES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> message_severity

SAP HANA DI Container-Specific Call**Available Parameters**

<code><container>#DI.GRANT_CONTAINER_API_PRIVILEGES_WITH_GRANT_OPTION</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> message_severity
<code><container>#DI.GRANT_CONTAINER_SCHEMA_PRIVILEGES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> message_severity
<code><container>#DI.GRANT_CONTAINER_SCHEMA_ROLES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> message_severity
<code><container>#DI.LIST</code>	container_lock_wait_timeout ignore_files ignore_folders recursive trace_context trace_level.<trace topic> message_severity
<code><container>#DI.LIST_CONFIGURED_LIBRARIES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> message_severity
<code><container>#DI.LIST_DEPLOYED</code>	container_lock_wait_timeout ignore_files ignore_folders recursive trace_context trace_level.<trace topic> message_severity
<code><container>#DI.MAKE</code>	container_lock_wait_timeout max_parallel_jobs optimized_redeploy simulate_make skip_unchanged_expansions trace_context trace_level.<trace topic> treat_warnings_as_errors undeploy_dependent_recursively enable_make_enforcer message_severity

SAP HANA DI Container-Specific Call**Available Parameters**

<code><container>#DI.MAKE_ASYNC</code>	container_lock_wait_timeout max_parallel_jobs optimized_redeploy simulate_make skip_unchanged_expansions trace_context trace_level.<trace topic> treat_warnings_as_errors undeploy_dependent_recursively enable_make_enforcer message_severity
<code><container>#DI.READ</code>	container_lock_wait_timeout ignore_files ignore_folders recursive trace_context trace_level.<trace topic> message_severity
<code><container>#DI.READ_DEPLOYED</code>	container_lock_wait_timeout ignore_files ignore_folders recursive trace_context trace_level.<trace topic> message_severity
<code><container>#DI.REVOKE_CONTAINER_API_PRIVILEGES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> message_severity
<code><container>#DI.REVOKE_CONTAINER_SCHEMA_PRIVILEGES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> message_severity
<code><container>#DI.REVOKE_GRANT_CONTAINER_SCHEMA_ROLES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> message_severity
<code><container>#DI.STATUS</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> message_severity

SAP HANA DI Container-Specific Call

Available Parameters

<container>#DI.WRITE

container_lock_wait_timeout
trace_context
trace_level.<trace topic>
message_severity

<container>#DI.CONFIGURE_CONTAINER

⚠ Caution

Deprecated since SAP HANA 1.0 SPS 12.

Example: Calling a Container-Specific Procedure with Parameters Set

📄 Sample Code

```
-- prepare path content table
create table MY_PATH_CONTENT like _SYS_DI.TT_FILESFOLDERS_CONTENT;
insert into MY_PATH_CONTENT (PATH, CONTENT) values ('mypath/', '');
insert into MY_PATH_CONTENT (PATH, CONTENT) values ('mypath/
myfile1.hdbtable', 'ROW TABLE MY_TABLE (X INTEGER)');
insert into MY_PATH_CONTENT (PATH, CONTENT) values ('mypath/.hdiconfig', '{
"file_suffixes" : { "hdbtable" : { "plugin_name" : "com.sap.hana.di.table",
"plugin_version" : "12.0.0" } } }');

-- prepare parameters table
create table MY_PARAMETERS like _SYS_DI.TT_PARAMETERS;
insert into MY_PARAMETERS (KEY, VALUE) values ('container_lock_wait_timeout',
'10');
-- call procedure
call MY_CONTAINER#DI.WRITE(MY_PATH_CONTENT, MY_PARAMETERS, ?, ?, ?);
```

Available SAP HANA DI Parameters

The following table describes the parameters available in SAP HANA DI and their possible values.

Parameter	Possible values	Description
container_lock_wait_timeou t	0...2,147,483,647	Specifies the time (in milliseconds) a container operation waits for a locking conflict to clear. The default value is the value of the corresponding SAP HANA DI configuration parameter <code>connection.container_default_transaction_lock_wait_timeout</code> . For more information, see <i>SAP HANA DI Configuration Parameters</i> .

Parameter	Possible values	Description
accept_risk_of_database_corruption_by_container_import	true, false	Indicates that the user knows that a container import could potentially corrupt the database. The default value is "false".
enable_make_enforcer	true, false	Terminate all external database connections blocking a make call. The default value is true.
export_container_schema_data	true, false	Indicates that a container export should also export the data of the container schema. If set to "false", only the schema structure will be exported. The default value is "true".
ignore_deployed	true, false	Indicates if existing files in the deployed file system are to be ignored when dropping a container. The default value is "false".
ignore_errors	true, false	Indicates if errors during an SAP HANA DI call should be ignored, that is, execute and commit as many internal operations as possible. Failing operations are reported to the user. The default value is "false".
ignore_files	true, false	Indicates if files are to be ignored in the output when reading files. The default value is "false".
ignore_folders	true, false	Indicates if folders are to be ignored in the output when reading files. The default value is "false".
ignore_non_existing_paths	true, false	Indicates if paths that do not exist should be ignored, for example, when deleting folders. The default value is "false".
ignore_work	true, false	Indicates if existing files in the work file system are to be ignored when dropping a container. The default value is "false".

Parameter	Possible values	Description
<code>max_parallel_jobs</code>	0 ... 2,147,483,647	Specifies the maximum number of parallel jobs for graph execution and artifact deployment. The default value is 8.
<code>move_containers_to_default_group</code>	true, false	Indicates if all containers in a group should be moved to the default group <code>_SYS_DI</code> first before dropping the container group. The default value is "false".
<code>optimized_redeploy</code>	true, false	Indicates if the optimized redeployment strategy should be used instead of the undeploy-deploy mechanism when re-deploying artifacts. The default value is "true".
<code>recursive</code>	true, false	Indicates if folders are to be read or deleted recursively. The default value is "false".
<code>simulate_make</code>	true, false	Indicates if the make should run only in simulation mode. The default value is "false".
<code>trace_context</code>	request, container	Indicates if, during an SAP HANA DI request, all traces for trace topics configured using the <code>trace_level_<trace topic></code> parameter are written to a separate trace file in addition to the DI server trace file. If set to "request", a new trace file is created for the request. For container operations, if set to "container", a trace file for the corresponding container is created or appended to. There is no default value.
<code>trace_level.<trace topic></code>	Fatal, Error, Warning, Info, Interface, Debug, InterfaceFull, DebugFull	Specifies the trace level of a specific trace topic. <code><trace topic></code> may be an arbitrary SAP HANA trace topic. There is no default value.

Parameter	Possible values	Description
<code>treat_warnings_as_errors</code>	true, false	<p>Indicates if warnings during an SAP HANA DI call should be treated as errors.</p> <p>The default value is "false".</p>
<code>undeploy</code>	true, false	<p>Indicates if, in case of a call to configure libraries, files corresponding to a library to be removed should also be undeployed.</p> <p>The default value is "false".</p>
<code>undeploy_dependent_recur_sively</code>	true, false	<p>Indicates that all dependent artifacts should be considered for an undeployment.</p> <p>The default value is "false".</p>
<code>variant</code>	"providers", "provides_and_requires", "impacted", "depends"	<p>Specifies the variant of the SAP HANA DI (HDI) container-specific procedure <code><container>#DI.GET_DEPENDENCIES</code> to be used; the following values are permitted:</p> <ul style="list-style-type: none"> • "providers" Returns the providing file for a given database object • "provides_and_requires" Returns the provided and required database objects for a given artifact • "impacted" Returns the files that depend on (are affected by) the specified files. • "depends" Returns the files that the specified files depend on. <p>The default value is "providers".</p>
<code>message_severity</code>	INFO, WARNING, ERROR	<p>Specifies the minimum severity of the messages to be returned by a SAP HANA DI procedure call.</p> <p>The default value is INFO (all messages are returned).</p>

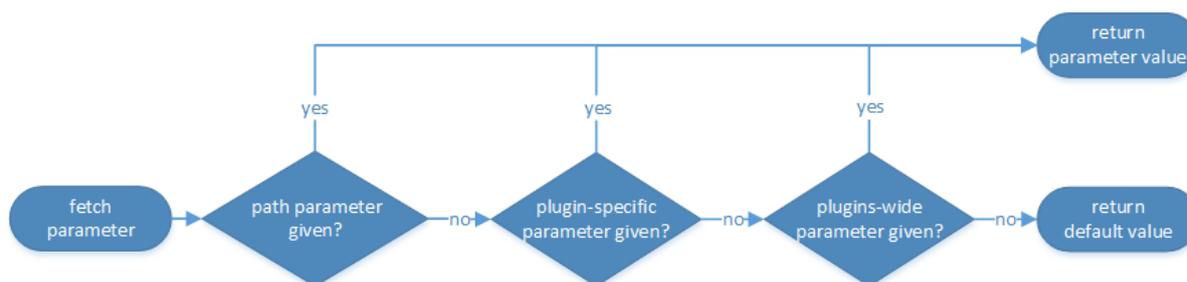
Parameters for Build Plugins

SAP HANA DI supports three types of parameters for controlling the execution flow of build plugins. On a global level, a “plugins-wide parameter” applies to all build plugins supporting the parameter. On a more fine-grained level, a “plugin-specific parameter” only applies to the specified build plugin. Eventually, a “path parameter” serves to control the handling of a single file.

The following table describes the structure of each parameter type.

Parameter Type	File	Parameter Structure
plugins-wide parameter	-	com.sap.hana.di/<key>
plugin-specific parameter	-	com.sap.hana.di.<plugin>/<key>
path parameter	<file>	<key>

From the point of view of a build plugin, the three types of parameters form a hierarchy whereby the most specific parameter type is considered first. The following diagram visualizes the process of fetching a parameter from the point of view of a build plugin.



The plugin-specific parameters support additional layering by allowing additional layers within the parameter structure. For example, if a requested parameter `com.sap.hana.di.<layer1>.<plugin1>/<key>` is not found, the build plugin automatically searches for a parameter `com.sap.hana.di.<layer1>/<key>`. The following table shows an example of layering of parameters.

Parameter Type	File	Parameter Structure
plugins-wide parameter	-	com.sap.hana.di/<key>
layered parameter	-	com.sap.hana.di.<layer1>/<key>
layered plugin parameter	-	com.sap.hana.di.<layer1>.<plugin1>/<key>
path parameter	<file>	<key>

The following section lists the available build plugin parameters and path parameters in SAP HANA DI. The section *Build Plugin Parameters* describes the available parameters for the configuration of build plugins.

Build Plugins

The following table lists the available parameters for the build plugins.

Build Plugin	Available Build Plugin Parameters	Available Path Parameters
Applies to all build plugins	optimized_redeploy skip_unchanged_expansions undeploy_dependent_recursively	-
com.sap.hana.di.cds	force_undeploy	-
com.sap.hana.di.role	force_undeploy	-
com.sap.hana.di.sequence	force_undeploy	-
com.sap.hana.di.table	force_undeploy	-
com.sap.hana.di.tabledata	batch_size	-

Example for calling the make procedure with a plugins-wide parameter set:

Sample Code

```
-- prepare deploy paths table
create table MY_DEPLOY_PATHS like _SYS_DI.TT_FILESFOLDERS;
insert into MY_DEPLOY_PATHS (PATH) values ('mypath/myfile1.hdbtable');
insert into MY_DEPLOY_PATHS (PATH) values ('mypath/.hdiconfig');
-- prepare parameters table with a plugins-wide parameter
create table MY_PARAMETERS like _SYS_DI.TT_PARAMETERS;
insert into MY_PARAMETERS (KEY, VALUE) values ('com.sap.hana.di /
force_undeploy', 'true');
-- call procedure
call MY_CONTAINER#DI.MAKE(MY_DEPLOY_PATHS, _SYS_DI.T_NO_FILESFOLDERS,
_SYS_DI.T_NO_FILESFOLDERS_PARAMETERS, MY_PARAMETERS, ?, ?, ?);
```

Example for calling the make procedure with a plugin-specific parameter set:

Sample Code

```
-- prepare deploy paths table
create table MY_DEPLOY_PATHS like _SYS_DI.TT_FILESFOLDERS;
insert into MY_DEPLOY_PATHS (PATH) values ('mypath/myfile1.hdbtable');
insert into MY_DEPLOY_PATHS (PATH) values ('mypath/.hdiconfig');
-- prepare parameters table with a plugin-specific parameter
create table MY_PARAMETERS like _SYS_DI.TT_PARAMETERS;
insert into MY_PARAMETERS (KEY, VALUE) values ('com.sap.hana.di.table/
force_undeploy', 'true');
-- call procedure
call MY_CONTAINER#DI.MAKE(MY_DEPLOY_PATHS, _SYS_DI.T_NO_FILESFOLDERS,
_SYS_DI.T_NO_FILESFOLDERS_PARAMETERS, MY_PARAMETERS, ?, ?, ?);
```

Example for calling the make procedure with a path parameter set:

Sample Code

```
-- prepare deploy paths table
create table MY_DEPLOY_PATHS like _SYS_DI.TT_FILESFOLDERS;
insert into MY_DEPLOY_PATHS (PATH) values ('mypath/myfile1.hdbtable');
insert into MY_DEPLOY_PATHS (PATH) values ('mypath/.hdiconfig');
-- prepare path parameters table
create table MY_PATH_PARAMETERS like _SYS_DI.TT_FILESFOLDERS_PARAMETERS;
insert into MY_PATH_PARAMETERS (PATH, KEY, VALUE) values ('mypath/
myfile1.hdbtable', 'force_undeploy', 'true');
-- call procedure
call MY_CONTAINER#DI.MAKE(MY_DEPLOY_PATHS, _SYS_DI.T_NO_FILESFOLDERS,
MY_PATH_PARAMETERS, _SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

Build Plugin Parameters

The following table describes the build plugin parameters available in SAP HANA DI and their possible values.

Build Plugin Parameter	Possible Values	Description
batch_size	0 ... 2,147,483,647	Specifies the batch size, for example, for batch database access or for batch processing within a build plugin. The default value is "1".
force_undeploy	true, false	Indicates if the undeployment of files should be forced within a build plugin that would alter an existing database object instead of simply re-creating it. The default value is "false".
optimized_redeploy	true, false	Indicates if the optimized redeployment strategy should be used instead of the undeploy-deploy mechanism when re-deploying artifacts. The default value is "true".
skip_unchanged_expansions	true, false	Instructs the make expander to not add expansions with same name and content as a deployed expansion to the deploy and undeploy sets. The default value is "false".

Build Plugin Parameter	Possible Values	Description
<code>undeploy_dependent_recursively</code>	true, false	Indicates that all dependent artifacts should be considered for an undeployment. The default value is "false".

Related Information

[SAP HANA DI Configuration Parameters \[page 1484\]](#)

11.5.4.2 SAP HANA DI Configuration Parameters

Configuration parameters are used to configure the behavior of SAP HANA DI. There are two types of configuration parameters: SAP HANA DI configuration parameters and container-specific configuration parameters.

SAP HANA DI configuration parameters configure the general behavior of SAP HANA DI. For example, they specify the time an SAP HANA DI operation waits for a locking conflict to clear or they specify the default behavior of containers.

Container-specific configuration parameters are used to control the behavior of a single container. For example, they specify the time a container operation waits for a locking conflict to clear or the maximum number of parallel jobs to be spawned during a make.

SAP HANA DI Configuration Parameters

Parameter	Possible Values	Description
<code>api.severity_for_invalid_parameter</code>	ERROR, WARNING, INFO	Specifies the severity of the corresponding log message when an invalid parameter has been passed with the SAP HANA DI operation. The default value is ERROR.
<code>blobs.container_default_days_to_keep</code>	-2,147,483,648 ... 2,147,483,647	Specifies the default number of days to keep data entries in the container-specific blob store. A value of 0 means deleting all entries. A negative value means keeping all entries. The default value is 10.

Parameter	Possible Values	Description
<code>blobs.transaction_lock_wait_timeout</code>	0 ... 2,147,483,647	<p>Specifies the time (in milliseconds) a blob store operation waits for a locking conflict to clear.</p> <p>The default value is 100,000.</p>
<code>connection.container_default_transaction_lock_wait_timeout</code>	0 ... 2,147,483,647	<p>Specifies the default time (in milliseconds) a container operation waits for a locking conflict to clear.</p> <p>The default value is the value of the corresponding SAP HANA DI configuration parameter <code>connection.global_transaction_lock_wait_timeout</code>.</p>
<code>connection.global_transaction_lock_wait_timeout</code>	0 ... 2,147,483,647	<p>Specifies the time (in milliseconds) an SAP HANA DI operation waits for a locking conflict to clear.</p> <p>The default value is 100,000.</p>
<code>connection.max_polls_for_master_indexserver</code>	0 ... 2,147,483,647	<p>Specifies the maximum number of polls for the master indexserver before aborting the SAP HANA DI operation.</p> <p>The default value is 100.</p>
<code>connection.poll_interval_for_master_indexserver</code>	0 ... 2,147,483,647	<p>Specifies the interval (in seconds) between polls for the master indexserver.</p> <p>The default value is 5.</p>
<code>connection.max_retries_for_initialization</code>	0 ... 2,147,483,647	<p>Specifies the maximum number of retries to initialize the DI-server during startup.</p> <p>The default value is 10.</p>
<code>make.default_max_parallel_jobs</code>	0 ... 2,147,483,647	<p>Specifies the default maximum number of parallel jobs to be spawned during a make.</p> <p>The default value is 4.</p>

Parameter	Possible Values	Description
<code>messages.container_default_days_to_keep</code>	-2,147,483,648 ... 2,147,483,647	<p>Specifies the default number of days to keep container-specific log messages. A value of 0 means all entries are deleted. A negative value means all entries are kept.</p> <p>The default value is the value of the corresponding SAP HANA DI configuration parameter <code>messages.global_days_to_keep</code>.</p>
<code>messages.container_default_requests_to_keep</code>	-2,147,483,648 ... 2,147,483,647	<p>Specifies the default number of requests to keep container-specific log messages. A value of 0 means all entries are deleted. A negative value means all entries are kept.</p> <p>The default value is the value of the corresponding SAP HANA DI configuration parameter <code>messages.global_requests_to_keep</code>.</p>
<code>messages.global_days_to_keep</code>	-2,147,483,648 ... 2,147,483,647	<p>Specifies the number of days to keep global SAP HANA DI log messages. A value of 0 means all entries are deleted. A negative value means all entries are kept.</p> <p>The default value is 10.</p>
<code>messages.global_requests_to_keep</code>	-2,147,483,648 ... 2,147,483,647	<p>Specifies the number of requests to keep global SAP HANA DI log messages. A value of 0 means all entries are deleted. A negative value means all entries are kept.</p> <p>The default value is 100.</p>
<code>messages.treat_errors_as_warnings</code>	A list of comma-separated error codes	<p>Specifies error codes whose corresponding messages should be treated as warnings instead of errors.</p> <p>The default value is an empty list.</p>
<code>trace.max_content_bytes_traced</code>	0 ... 2,147,483,647	<p>Specifies the maximum length (in bytes) of a content to be traced.</p> <p>The default value is 100.</p>

Example: Configuring SAP HANA DI with an SAP HANA DI Configuration Parameter

Sample Code

```
-- prepare configuration parameters table
create table MY_CONFIG_PARAMETERS like _SYS_DI.TT_PARAMETERS;
insert into MY_CONFIG_PARAMETERS(KEY, VALUE) values
('make.default_max_parallel_jobs', '10');
-- prepare parameters table
create table MY_PARAMETERS like _SYS_DI.TT_PARAMETERS;
-- call procedure
call _SYS_DI.CONFIGURE_DI_PARAMETERS(MY_CONFIG_PARAMETERS,
MY_PARAMETERS, ?, ?, ?);
```

Container-Specific Configuration Parameters

The following table describes the container-specific configuration parameters and their possible values.

SAP HANA DI Container-specific Configuration Parameters

Parameter	Possible Values	Description
<code>blobs.days_to_keep</code>	-2,147,483,648 ... 2,147,483,647	Specifies the number of days to keep data entries in the blob store. A value of 0 means all entries are deleted. A negative value means all entries are kept. The default value is the value of the corresponding SAP HANA DI configuration parameter <code>blobs.container_default_days_to_keep</code> .
<code>connection.transaction_lock_wait_timeout</code>	0 ... 2,147,483,647	Specifies the time (in milliseconds) a container operation waits for a locking conflict to clear. The default value is the value of the corresponding SAP HANA DI configuration parameter <code>connection.container_default_transaction_lock_wait_timeout</code> .
<code>make.force_logical_schema_targets</code>	"true" or "false"	Indicates if the usage of logical schemas in design-time artifacts should be forced. The default value is "false".

Parameter	Possible Values	Description
<code>make.max_parallel_jobs</code>	0 ... 2,147,483,647	<p>Specifies the maximum number of parallel jobs to be spawned during a make.</p> <p>The default value is the value of the corresponding SAP HANA DI configuration parameter <code>make.default_max_parallel_jobs</code>.</p>
<code>make.prohibit_config_file</code>	"true" or "false"	<p>Indicates if the deployment of configuration files (<code>.hdi.config</code>) should be disabled.</p> <p>The default value is "false".</p>
<code>make.prohibit_definer_mode</code>	"true" or "false"	<p>Indicates if the creation of "definer-mode" procedures is disabled. The default value is "false".</p>
<code>make.prohibit_namespace_file</code>	"true" or "false"	<p>Indicates if the deployment of namespace files (<code>.hdi.namespace</code>) should be disabled.</p> <p>The default value is "false".</p>
<code>make.prohibit_table_creation</code>	"true" or "false"	<p>Indicates if the creation of tables is disabled. The default value is "false".</p>
<code>messages.days_to_keep</code>	-2,147,483,648 ... 2,147,483,647	<p>Specifies the number of days to keep log messages. A value of 0 means all entries are deleted. A negative value means all entries are kept.</p> <p>The default value is the value of the corresponding SAP HANA DI configuration parameter <code>messages.container_default_days_to_keep</code>.</p>
<code>messages.requests_to_keep</code>	-2,147,483,648 ... 2,147,483,647	<p>Specifies the number of requests to keep log messages. A value of 0 means all entries are deleted. A negative value means all entries are kept.</p> <p>The default value is the value of the corresponding SAP HANA DI configuration parameter <code>messages.container_default_requests_to_keep</code>.</p>

Parameter	Possible Values	Description
build.plugin.name/disabled	"true" or "false"	Indicated if the build plugin specified by the parameter is disabled. By default, there is no such parameter set; no build plugin is disabled by default..

Example: Configuring a Container with a Container-specific Configuration Parameter

Sample Code

```
-- prepare configuration parameters table
create table MY_CONFIG_PARAMETERS like _SYS_DI.TT_PARAMETERS;
insert into MY_CONFIG_PARAMETERS(KEY, VALUE) values
('make.max_parallel_jobs', '10');
-- prepare parameters table
create table MY_PARAMETERS like _SYS_DI.TT_PARAMETERS;
-- call procedure
call MY_CONTAINER#DI.CONFIGURE_CONTAINER_PARAMETERS(MY_CONFIG_PARAMETERS,
MY_PARAMETERS, ?, ?, ?);
```

11.5.5 List Plug-in Libraries That Can Be Configured for a Container

It may be useful to find out which HDI plug-in libraries and versions are available in the database and can be configured for use in a container.

Procedure

1. In an SQL console, connect to the database as the HDI administrator.
2. Display a list of plug-in libraries available in the HDI container.

Run the following SQL statement

```
CALL _SYS_DI.LIST_LIBRARIES(_SYS_DI.T_NO_PARAMETERS, ?, ?, ?, ?);
```

Confirm that the SQL code completes successfully and displays the HDI return code 0.

3. Examine the result set returned by SQL.

The following example shows an excerpt of a typical result set for this request:

Output Code

Excerpt of result set

```
LIBRARY_NAME;LIBRARY_VERSION;PLUGIN_ID;PLUGIN_VERSION
com.sap.hana.di.afllangprocedure;0.0;com.sap.hana.di.afllangprocedure;
2.0.10.0
com.sap.hana.di.analyticprivilege;0.0;com.sap.hana.di.analyticprivilege;
2.0.10.0
com.sap.hana.di.calculationview;0.0;com.sap.hana.di.calculationview;
2.0.10.0
com.sap.hana.di.cds;0.3;com.sap.hana.di.cds;2.0.10.0
com.sap.hana.di.constraint;0.0;com.sap.hana.di.constraint;2.0.10.0
com.sap.hana.di.copyonly;0.0;com.sap.hana.di.copyonly;2.0.10.0
com.sap.hana.di.dropcreatetable;0.0;com.sap.hana.di.dropcreatetable;
2.0.10.0
com.sap.hana.di.flowgraph;0.0;com.sap.hana.di.flowgraph;2.0.10.0
```

11.5.6 Create a Container Group

A container group is used for administrating a set of containers. Every container group can be managed by different users.

Procedure

1. In an SQL console, connect to the database as the HDI administrator.
2. Insert the following SQL statement:

```
CALL _SYS_DI.CREATE_CONTAINER_GROUP('<container_group_name>',
  _SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

Note

Replace the name of the container group `<container_group_name>` in the SQL `CALL` statement with your container-group name

3. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

4. Confirm that the new container group has been created.

Tip

The system view `_SYS_DI.M_ALL_CONTAINER_GROUPS` contains a list of **all** created HDI container groups.

11.5.7 Drop a Container Group

The HDI administrator can drop a container group.

Prerequisites

The container group is empty.

Procedure

1. In an SQL console, connect to the database as the HDI administrator.
2. Drop the specified container group.

Insert the following SQL statement into the SQL console:

```
CALL _SYS_DI.DROP_CONTAINER_GROUP('<container_group_name>',  
_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

i Note

Replace the name of the container group `<container_group_name>` in the SQL `CALL` statement with your container-group name

3. Execute the SQL code.
Confirm that the SQL code completes successfully and displays the HDI return code 0.
4. Confirm that the schema `_SYS_DI#<container_group_name>` no longer exists.

11.5.8 Grant Container Group Administrator Privileges to Another User

Every container group can have its own set of administrators. An HDI administrator or a container group administrator must explicitly grant administrative privileges on a container group.

Context

An HDI administrator can grant another user container group administrator privileges to any container group, whereas a container group administrator can only grant another user the container group administrator privileges for his own container group. This method uses the predefined `_SYS_DI.T_DEFAULT_CONTAINER_GROUP_ADMIN_PRIVILEGES` table, which contains the largest possible

set of privileges that can be granted for a user of this type. You can reduce the set of privileges granted by explicitly specifying the desired set of privileges and not using this default table.

Note

A variant of this procedure,

`_SYS_DI.GRANT_CONTAINER_GROUP_API_PRIVILEGES_WITH_GRANT_OPTION`, also exists; it grants the specified privileges `WITH GRANT OPTION` to the target user. This variant is only needed in special scenarios, for example, when building a custom SQL API by wrapping the HDI SQL API in SQLScript procedures.

Procedure

1. In an SQL console, connect to the database as the HDI administrator.
2. Grant container-group administrator privileges to a specified user.

Insert the following SQL statement into the SQL console:

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
_SYS_DI.TT_API_PRIVILEGES;
INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME) SELECT
'<new_container_group_admin_username>', PRIVILEGE_NAME, OBJECT_NAME FROM
_SYS_DI.T_DEFAULT_CONTAINER_GROUP_ADMIN_PRIVILEGES;
CALL _SYS_DI.GRANT_CONTAINER_GROUP_API_PRIVILEGES('<container_group_name>',
#PRIVILEGES, _SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;
```

→ Tip

Replace the name of the user `<container_group_admin_username>` in `INSERT` command in line 2 with the name of the user to whom the API privileges should be granted, and replace the name of the container group `<container_group_name>` in the `CALL` command in line 3 with the name of the desired container group name.

3. Execute the SQL code.
Confirm that the SQL code completes successfully and displays the HDI return code 0.
4. Confirm that the new container group administrator is able to call HDI API procedures in the API schema `_SYS_DI#<container_group_name>`.

11.5.9 Revoke Container Group Administrator Privileges from a Container Group Administrator

An HDI administrator can revoke administration privileges on any container group from a user.

Procedure

1. In an SQL console, connect to the database as the HDI administrator.
2. Revoke the specified container-group administrator privileges.

Insert the following SQL statement into the SQL console:

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
SYS_DI.TT_API_PRIVILEGES;
INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME) SELECT
'NEW_CONTAINER_GROUP_ADMIN', PRIVILEGE_NAME, OBJECT_NAME FROM
SYS_DI.T_DEFAULT_CONTAINER_GROUP_ADMIN_PRIVILEGES WHERE NOT (PRIVILEGE_NAME
= 'SELECT' AND OBJECT_NAME LIKE 'SYS_DI.T%');
CALL SYS_DI.REVOKE_CONTAINER_GROUP_API_PRIVILEGES('G', #PRIVILEGES,
SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;
```

→ Tip

Replace the name of the user `<NEW_CONTAINER_GROUP_ADMIN>` in `INSERT` command in line 2 with the name of the user from whom the API privileges should be revoked, and replace the name of the container group `<G>` in the `CALL` command in line 3 with the name of the desired container group name.

3. Execute the SQL code.
Confirm that the SQL code completes successfully and displays the HDI return code 0.
4. Confirm that the container group administrator is no longer able to call HDI API procedures in the API schema `_SYS_DI#<container_group_name>`.

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

[Grant Container-Group Administrator Privileges to a User \[page 1496\]](#)

11.5.10 Move a Container to Another Container Group

The HDI administrator can move a container to another container group.

Procedure

1. In an SQL console, connect to the database as the HDI administrator.
2. Move the specified container to the new target container group.

Insert the following SQL statement into the SQL console:

```
CALL _SYS_DI.MOVE_CONTAINER_TO_GROUP('<container_group_to_be_moved>',  
'<target_container_group>',_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

3. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

4. Using the `M_ALL_CONTAINERS` monitoring view, check that `<container_group_to_be_moved>` has the container group `<target_container_group>` assigned in the `CONTAINER_GROUP_NAME` column.

11.6 Maintaining HDI Container Groups

An HDI container-group administrator manages sets of containers collected in one or more container groups assigned by the HDI administrator.

Managing an HDI container group typically involves the following administrator tasks:

- Grant and revoke container-group administrator privileges
- Create and drop containers
- Grant and revoke container administrator privileges
- Grant and revoke container access

→ Tip

The APIs of a container group named “G” are in the schema `_SYS_DI#G`.

HDI Container Administration

The HDI container group administrator can perform the same administrative tasks as an HDI container administrator. For details about the functionality available to the HDI container administrator, see the section about HDI container administration in *Related Information* below. To perform container-related administration tasks, the container-group administrator calls the appropriate HDI container administration SQL procedures, not of the target container, but of the container **group schema** (for example, `_SYS_DI#G`) of the HDI container

group administrator. The name of the target container is specified by means of an additional first parameter, as illustrated in the sample below.

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
_SYS_DI.TT_SCHEMA_PRIVILEGES;
INSERT INTO #PRIVILEGES ( PRIVILEGE_NAME, PRINCIPAL_SCHEMA_NAME,
PRINCIPAL_NAME ) VALUES ( 'SELECT', '', 'U' );
CALL _SYS_DI#G.GRANT_CONTAINER_SCHEMA_PRIVILEGES( 'C', #PRIVILEGES,
_SYS_DI.TT_NO_PARAMETERS, ?, ?, ? );
DROP TABLE #PRIVILEGES;
```

The example above shows how to grant a user “U” the privileges to access the run-time objects in the container “C”. The HDI container group administrator calls the SQL procedure `GRANT_CONTAINER_SCHEMA_PRIVILEGES` in the `_SYS_DI` schema, passing the container’s name “C” as the additional first parameter:

Note

Container-administration tasks should normally be performed by the container’s assigned administrator. The HDI container **group** administrator should only be used to perform every-day container administration tasks in exceptional circumstances. The only exception to this rule is the creation and dropping of containers, which can only be performed by the container **group** administrator.

Exporting and Importing Containers

In special cases, a container and its dependencies can be exported from the database and imported into a different database, for, example, when a container needs to be copied from one container to another. For more information about how to export and import containers for copy purposes, see *Related Information* below.

It is also possible to export and import a container for support purposes. However, it is not recommended and is not without risk, as described in the following caution note.

Caution

The API procedures `_SYS_DI#G.EXPORT_CONTAINER_FOR_SUPPORT` and `_SYS_DI#G.IMPORT_CONTAINER_FOR_SUPPORT` are available to the HDI container-group administrator but intended for use by SAP support, exclusively. The exported data might include private or confidential data from the container, and the container import could also compromise the integrity of the database.

Related Information

[Maintaining HDI Containers \[page 1508\]](#)

[Maintaining the HDI \[page 1465\]](#)

[SAP HANA Deployment Infrastructure \[page 1450\]](#)

[Export a Container for Copy Purposes \[page 1505\]](#)

11.6.1 Grant Container-Group Administrator Privileges to a User

Container-group administrator privileges can be granted to another user at any time.

Context

Each container group can have its own set of administrators. Administrative privileges for a container group must be explicitly granted by an HDI administrator, or the container group's administrator. Unlike the HDI administrator, the container-group administrator can only grant another user the container group administrator privileges for their own container groups.

This method uses the predefined `_SYS_DI.T_DEFAULT_CONTAINER_GROUP_ADMIN_PRIVILEGES` table, which contains the largest possible set of privileges that can be granted for a user of this type. However, it is possible to reduce the set of privileges granted by specifying the desired set of privileges explicitly and not using this default table.

→ Tip

The procedure `_SYS_DI#G.GRANT_CONTAINER_GROUP_API_PRIVILEGES_WITH_GRANT_OPTION` can also be used; it grants the given privileges "WITH GRANT OPTION" to the target user. This procedure is only needed in special scenarios, for example, when building a custom SQL API by wrapping the HDI SQL API in SQLScript procedures.

Procedure

1. In an SQL console, connect to the database with an HDI administrator user.
2. Open the SQL editor for this database.
3. Insert the following SQL code into the SQL editor:

≡ Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
_SYS_DI.TT_API_PRIVILEGES;
INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME)
SELECT 'OTHER_CONTAINER_GROUP_ADMIN', PRIVILEGE_NAME, OBJECT_NAME FROM
_SYS_DI.T_DEFAULT_CONTAINER_GROUP_ADMIN_PRIVILEGES;
CALL _SYS_DI#G.GRANT_CONTAINER_GROUP_API_PRIVILEGES(#PRIVILEGES,
_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;
```

4. Adjust the name of the user `OTHER_CONTAINER_GROUP_ADMIN` in the `INSERT` command in line 2 to reflect the name of the user to whom the API privileges should be granted, and the name of the container group "G" in line 3 to correspond with the name of your container group.
5. Execute the SQL code.

- (Optional) Confirm that the `OTHER_CONTAINER_GROUP_ADMIN` user is now able to call HDI API procedures in the container group "G"'s API schema `_SYS_DI#G`, where "G" is replaced with the name of your HDI container group.

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

[Revoke Container-Group Administrator Privileges from an Administrator User \[page 1497\]](#)

11.6.2 Revoke Container-Group Administrator Privileges from an Administrator User

Container group administration privileges can be revoked from a user at any time.

Context

Each container group can have its own set of administrators. Administrative privileges for a container group must be explicitly granted and revoked by an HDI administrator, or a container-group's own administrator. Unlike the HDI administrator, the container-group administrator can only revoke container-group administrator privileges from another user for their own container groups.

→ Tip

This method uses the predefined table `_SYS_DI.T_DEFAULT_CONTAINER_GROUP_ADMIN_PRIVILEGES`, which contains the largest possible set of privileges that can be granted for a user of this type. However, it is possible to reduce the set of privileges granted by specifying the desired set of privileges individually and explicitly instead.

Procedure

- In an SQL console, connect to the database with an HDI administrator user.
- Open the SQL editor for this database.
- Revoke the container-group administrator privileges.

Insert the following SQL code into the SQL editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
  _SYS_DI.TT_API_PRIVILEGES;
```

```

INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME)
SELECT 'OTHER_CONTAINER_GROUP_ADMIN', PRIVILEGE_NAME, OBJECT_NAME FROM
SYS_DI.T_DEFAULT_CONTAINER_GROUP_ADMIN_PRIVILEGES WHERE NOT
(PRIVILEGE_NAME = 'SELECT' AND OBJECT_NAME LIKE 'SYS_DI.T%');
CALL SYS_DI#G.REVOKE_CONTAINER_GROUP_API_PRIVILEGES(#PRIVILEGES,
SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;

```

- a. Adjust the name of the user `OTHER_CONTAINER_GROUP_ADMIN` in the `INSERT` statement in line 2 to reflect the name of the user from whom the API privileges should be revoked, and the name of the container group “G” in the `CALL` statement in line 3 to correspond with the name of your container group.
4. Execute the SQL code.

Check that the code completes successfully with the HDI return code 0.
 5. (Optional) Confirm that the `OTHER_CONTAINER_GROUP_ADMIN` user is no longer able to call HDI API procedures in the container group “G”'s API schema `SYS_DI#G`, where “G” is replaced with the name of your HDI container group.

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

[Grant Container-Group Administrator Privileges to a User \[page 1496\]](#)

11.6.3 Create a Container

The HDI container-group administrator can create a new container.

Context

HDI container-group administrators can create HDI containers in any HDI container group for which they are responsible.

Procedure

1. In an SQL console, connect to the database with an HDI administrator user.
2. Open the SQL editor for this database.
3. Create the container.

Insert the following SQL code into the SQL editor:

Sample Code

```
CALL _SYS_DI#G.CREATE_CONTAINER('C', _SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

4. Adjust the name of the container group "G" and the name of the new container "C" to suit the names of your container group and container respectively.
5. Execute the SQL code.

Check that the code completes successfully with the HDI return code 0.

6. (Optional) Confirm that the container exists by checking that a corresponding new entry in the container group schema's `_SYS_DI#G.M_CONTAINERS` view is now present.

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

[Drop a Container \[page 1499\]](#)

11.6.4 Drop a Container

The HDI container-group administrator can drop a container.

Context

HDI container-group administrators can drop HDI containers from any HDI container group for which they are responsible.

Procedure

1. In an SQL console, connect to the database with an HDI administrator user.
2. Open the SQL editor for this database.
3. Drop the container.

Insert the following SQL code into the SQL editor:

Sample Code

```
CALL _SYS_DI#G.DROP_CONTAINER('C', _SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

4. Adjust the name of the container group "G" and the name of the new container "C" to suit the names of your container group and container respectively.

5. Execute the SQL code.

Check that the code completes successfully with the HDI return code 0.

6. (Optional) Confirm that the container no longer exists by checking that a corresponding entry in the container group schema's `_SYS_DI#G.M_CONTAINERS` view is not present anymore.

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

[Create a Container \[page 1498\]](#)

11.6.5 Grant Container Administrator Privileges to a User

Container administrator privileges can be granted to another user at any time.

Context

Each container can have its own set of administrators. Administrative privileges for a container must be explicitly granted or revoked either by an HDI container group administrator or an HDI container administrator with the necessary privileges.

i Note

This method uses the predefined table `_SYS_DI.T_DEFAULT_CONTAINER_ADMIN_PRIVILEGES`, which contains the largest possible set of privileges that can be granted for a user of this type. You can reduce the set of privileges granted to a user by explicitly specifying the desired set of privileges and not using this default table.

Procedure

1. In an SQL console, connect to the database as the administrator of the target HDI container group (for example, "G").
2. Open the SQL editor for this database.
3. Grant container-administrator privileges.

Insert the following SQL code into the editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
_SYS_DI.TT_API_PRIVILEGES;
```

```
INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME)
SELECT 'NEW_CONTAINER_ADMIN', PRIVILEGE_NAME, OBJECT_NAME FROM
SYS_DI.T_DEFAULT_CONTAINER_ADMIN_PRIVILEGES;
CALL SYS_DI#G.GRANT_CONTAINER_API_PRIVILEGES('C', #PRIVILEGES,
SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;
```

4. Adjust the name of the user `NEW_CONTAINER_ADMIN` in the `INSERT` statement in line 2 to reflect the name of the user to whom the API privileges should be granted. You will also have to change the name of the container group “G” in the API schema `_SYS_DI#G` and the name of the container “C” to reflect the names in your landscape.
5. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.
6. (Optional) Confirm that the `NEW_CONTAINER_ADMIN` user is now able to call HDI container API procedures in the container C's API schema `C#DI`.

Related Information

[Revoke Container Administrator Privileges from a User \[page 1501\]](#)

[Maintaining HDI Container Groups \[page 1494\]](#)

11.6.6 Revoke Container Administrator Privileges from a User

Container administrator privileges can be granted to another user at any time.

Context

Each container can have its own set of administrators. Administrative privileges for a container must be explicitly granted or revoked by an HDI container group administrator, or an HDI container administrator with the necessary privileges.

i Note

This method uses the predefined table `_SYS_DI.T_DEFAULT_CONTAINER_ADMIN_PRIVILEGES`, which contains the largest possible set of privileges that can be granted for a user of this type. You can reduce the set of privileges granted to a user by explicitly specifying the desired set of privileges and not using this default table.

Procedure

1. In an SQL console, connect to the database as the administrator of the HDI container group "G".
2. Open the SQL editor for this database.
3. Insert the following SQL code into the editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
_SYS_DI.TT_API_PRIVILEGES;
INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME)
SELECT 'NEW_CONTAINER_ADMIN', PRIVILEGE_NAME, OBJECT_NAME FROM
_SYS_DI.T.DEFAULT_CONTAINER_ADMIN_PRIVILEGES WHERE NOT (PRIVILEGE_NAME =
'SELECT' AND OBJECT_NAME LIKE '_SYS_DI.T%');
CALL _SYS_DI#G.REVOKE_CONTAINER_API_PRIVILEGES('C', #PRIVILEGES,
_SYS_DI.T.NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;
```

4. Adjust the name of the user `NEW_CONTAINER_ADMIN` in the `INSERT` statement in line 2 to reflect the name of the user from whom the API privileges should be revoked. In the `CALL` statement in line 3, you will also have to change the name of the container group "G" in the API schema `_SYS_DI#G` and the name of the container "C" to reflect the names in your landscape.
5. Execute the SQL code.
Confirm that the SQL code completes successfully and displays the HDI return code 0.
6. (Optional) Confirm that the `NEW_CONTAINER_ADMIN` user is no longer able to call HDI container API procedures in the container C's API schema `C#DI`.

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

[Grant Container Administrator Privileges to a User \[page 1500\]](#)

11.6.7 Grant a Support User Access to a Container

Provide members of the support teams with temporary access to a container.

Context

In the event of container-related problems, it might be necessary for a support user to access HDI-internal objects in a container API schema, for example, `C#DI` for the container "C". These privileges must be temporarily granted to an explicit support user and then revoked when the support task is completed.

⚠ Caution

The use of this function is only recommended in exceptional circumstances; it allows the support user to access all the data in the container, some of which could be private or confidential. Use of this function could also compromise the integrity of the container resulting in an unusable container or data loss. Enabling a support user to access a container raises a security alert to the database administrator.

Procedure

1. In an SQL console, connect to the database with the administrator of the HDI container group G.
2. Open the SQL editor for this database.
3. Insert the following SQL code into the editor:

⌘ Sample Code

```
CALL _SYS_DI#G.GRANT_CONTAINER_SUPPORT_PRIVILEGE('C', 'SELECT',  
'CONTAINER_SUPPORT_USER', _SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

4. Adjust the schema name of the container group G's API schema `_SYS_DI#G`.
5. Adjust the name of the container "C".
6. Adjust the privilege as needed.

Possible values are `SELECT`, `UPDATE`, `INSERT`, and `DELETE`.
7. Adjust the name of the user `CONTAINER_SUPPORT_USER`.
8. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.
9. (Optional) Confirm that the `CONTAINER_SUPPORT_USER` user can now access objects in the container C's API schema `C#DI` per the given privileges.

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

[Revoke Access to a Container from a Support User \[page 1504\]](#)

11.6.8 Revoke Access to a Container from a Support User

Revoke privileges granted to members of the support teams for temporary access to containers.

Context

In the event of container-related problems, it might be necessary for a support user to access HDI-internal objects in a container API schema, for example, C#DI for the container "C". These privileges must be temporarily granted to an explicit support user and then revoked when the support task is completed.

Procedure

1. In an SQL console, connect to the database with the administrator of the HDI container group G.
2. Open the SQL editor for this database.
3. Insert the following SQL code into the editor:

Sample Code

```
CALL _SYS_DI#G.REVOKE_CONTAINER_SUPPORT_PRIVILEGE( 'C', 'SELECT',  
'CONTAINER_SUPPORT_USER', _SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

- a. Adjust the schema name of the container group G's API schema `_SYS_DI#G`.
 - b. Adjust the name of the container "C".
 - c. Adjust the privilege as needed.

Possible values are `SELECT`, `UPDATE`, `INSERT`, and `DELETE`.
 - d. Adjust the name of the user `CONTAINER_SUPPORT_USER`.
4. Execute the SQL code.
 5. (Optional) Confirm that the `CONTAINER_SUPPORT_USER` user can no longer access objects in the container C's API schema `C#DI` per the given privileges.

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

[Grant a Support User Access to a Container \[page 1502\]](#)

11.6.9 Export a Container for Copy Purposes

A container and its dependencies can be exported to a table or a file, which can then be used to import the container into another database.

Prerequisites

In this task, the example used assumes the following:

- Container group `G` exists
- Container `C` is assigned to container group `G`

i Note

Adjust the names to suit your local requirements.

Context

The export of a source container, for example, `C1`, from a container group called `G` is performed by calling the built-in procedure `_SYS_DI#G.EXPORT_CONTAINER_FOR_COPY`. The procedure expects as inputs the name of the source container, the schema, and the names of two tables: the table into which the export data will be written (`EXPORT_TABLE`), and a table containing any parameters. After the export procedure has completed successfully, the specified table contains not only the objects of the source container's API schema (`C1#DI`) but also the deployed objects of the source container's run-time schema (`C1`), including all dependent data.

To export a container `C1` in container group `G` for copy purposes, perform the following steps:

Procedure

1. Open an SQL console and connect to the SAP HANA database with the permissions of the administrator of the HDI container group `G`.
2. Open the SQL editor for this database.
3. Paste the following SQL code into the SQL editor.

```
CREATE COLUMN TABLE EXPORT_TABLE LIKE _SYS_DI.TT_CONTAINER_EXPORT;  
GRANT INSERT ON EXPORT_TABLE TO C1#DI;  
GRANT SELECT ON EXPORT_TABLE TO C1#DI;  
CALL _SYS_DI#G.EXPORT_CONTAINER_FOR_COPY  
('C1', CURRENT_SCHEMA, 'EXPORT_TABLE', _SYS_DI.T_NO_PARAMETERS, ?, ?, ?);  
EXPORT EXPORT_TABLE AS BINARY INTO '<export_path>'  
WITH REPLACE NO DEPENDENCIES;
```

- a. If necessary, change the name of the container group `G` and the container `C1` to suit the names of the source objects in your environment.

i Note

If the exported container is copied within the same database, the `EXPORT...` statement can be skipped. If the target is another database, the `EXPORT...` statement creates a file on the file system that can be copied to the file system in the target database.

- b. If necessary, change the name of the target table ("`EXPORT_TABLE`" in the example above) to the name of the table that should receive the exported container.

! Restriction

The target table cannot be a temporary table.

4. Execute the SQL code, and check that the operation completed successfully (HDI return code 0).
5. Confirm that the container has been exported to the `EXPORT_TABLE` (or to the specified target file).

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

[Import a Container for Copy Purposes \[page 1506\]](#)

11.6.10 Import a Container for Copy Purposes

A container and its dependencies can be imported into the same (or another) database from a table or file.

Context

The import of a source container, for example, `C1` into a new container `C2` is performed by calling the built-in procedure `_SYS_DI#G.IMPORT_CONTAINER_FOR_COPY`. This procedure expects as input the name of an existing empty target container `C2`, the schema and table names of the table containing the exported results, and a table containing any necessary parameters. The import procedure writes the data from the table `EXPORT_TABLE` to the target container, runs a `make`, and copies back the existing data of the source container's run-time schema contained in the export table. After the import procedure has completed successfully, the target container will be a copy of the source container including all the data from the original run-time schema.

i Note

Only objects deployed by the source container's `MAKE` procedure are copied. Objects created manually in the source container's run-time schema are not copied.

To import a container C1 into a new container C2 for copy purposes, perform the following steps:

! Restriction

The container version of the target container must be the same as the container version of the source container. You can check the container version with the view `_SYS_DI#G.M_CONTAINER_VERSIONS`.

Procedure

1. Open an SQL console and connect to the SAP HANA database with the permissions of the administrator of the HDI container group G.
2. Open the SQL editor for this database.
3. If needed, create the new target container, for example, C2.

```
CALL _SYS_DI#G.CREATE_CONTAINER('C2', _SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

i Note

The new container must be empty for the purpose of an import.

4. If necessary, grant the object owner of the target container C2#00 any required privileges for all the external objects required by the container to be imported.

i Note

All external objects referenced by the container's design-time objects must be available and accessible to the object owner of the target container.

5. Paste the following SQL code into the SQL editor.

```
IMPORT EXPORT_TABLE FROM '<export_path>';  
CALL _SYS_DI#G.IMPORT_CONTAINER_FOR_COPY  
('C2', CURRENT_SCHEMA, 'EXPORT_TABLE', _SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

i Note

If the imported container is copied from within the same database, the `IMPORT...` statement can be skipped.

6. Execute the SQL code, and check that the operation completed successfully (HDI return code 0).
7. Confirm that the container has been successfully imported.

In this example, you can check that the expected run-time objects of source container C1 are also present in the target container C2.

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

11.7 Maintaining HDI Containers

An HDI container administrator configures and controls access to a container.

The SAP HANA Deployment Infrastructure (HDI) provides a service that enables you to deploy database development artifacts to so-called containers. This service includes a family of consistent design-time artifacts for all key SAP HANA platform database features which describe the target (run-time) state of SAP HANA database artifacts, for example: tables, views, or procedures. These artifacts are modeled, staged (uploaded), built, and deployed into SAP HANA.

The SAP HANA service broker is used to create and destroy HDI containers; each HDI container comprises a design-time container (DTC), which is an isolated environment used to store design-time files, and a run-time container (RTC), which is used to store deployed objects built according to the specification stored in the corresponding design-time artifacts.

The deployment process populates the database run-time with the specified catalog objects. In addition to database artifacts, HDI also enables you to import and export table content such as business configuration data and translatable texts.

! Restriction

HDI enables you to deploy database objects only; it is not possible (or necessary) to deploy application-layer artifacts such as JavaScript programs or OData objects.

The HDI container administrator manages one or more containers assigned by the container-group administrator. The role of the container-manager focuses primarily on configuring and controlling access to the HDI containers used to store the database objects deployed by the SAP HANA Deployment Infrastructure deploy service and repairing any problems that occur with run-time objects in the assigned HDI containers. An HDI container administrator can manage one or more containers in one HDI container group or multiple containers distributed across multiple container groups.

The configuration of HDI containers also involves the creation and configuration of the following design-time artifacts:

- Container deployment configuration (`.hdiconfig`)
A JSON file containing a list of the bindings between database artifact types (for example, sequence, procedure, table) and the corresponding deployment plug-in (and version).
- Run-time container namespace rules (`.hdinamespace`)
A JSON file containing a list of design-time file suffixes and the naming rules for the corresponding run-time locations.

→ Tip

The APIs of a container named "C" are in the schema `C#DI`.

HDI-Container Maintenance Tasks

To manage an HDI container group, the administrator performs the following common tasks:

- Grant and revoke container administrator privileges
- Grant and revoke access to the container development API
- Grant and revoke access to a container's schema
- Grant and revoke a user role from the container's schema
- List all currently configured build plug-in libraries available to a container
- Configure the default set of build plug-in libraries available to a container
- Configure a custom set of build plug-in libraries available to a container
- Configure container parameters

Related Information

[Maintaining HDI Container Groups \[page 1494\]](#)

[Maintaining the HDI \[page 1465\]](#)

[SAP HANA Deployment Infrastructure \[page 1450\]](#)

11.7.1 Grant HDI Container Administrator Privileges to a User

Enable administrator access to an HDI container.

Context

HDI container administrator privileges are initially granted to a user by an administrator of the container group that the container belongs to. If these privileges have been granted "with grant option", the HDI container administrator can also grant these privileges to another user as described below:

i Note

The predefined table `_SYS_DI.T_DEFAULT_CONTAINER_ADMIN_PRIVILEGES` used in this task contains the largest possible set of privileges that can be granted for a user of this type. You can reduce the set of privileges granted by explicitly specifying the desired set of privileges and not using this default table.

Procedure

1. In an SQL console, connect to the database with an administrator of the HDI container "C".

2. Open the SQL editor for this database.
3. Insert the following SQL code into the editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
_SYS_DI.TT_API_PRIVILEGES;
INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME)
SELECT 'NEW_CONTAINER_ADMIN', PRIVILEGE_NAME, OBJECT_NAME FROM
_SYS_DI.T_DEFAULT_CONTAINER_ADMIN_PRIVILEGES;
CALL C#DI.GRANT_CONTAINER_API_PRIVILEGES(#PRIVILEGES,
_SYS_DI.T_NO_PARAMETERS, '?', '?', '?');
DROP TABLE #PRIVILEGES;
```

- a. Adjust the name of the user "NEW_CONTAINER_ADMIN" in line 2 to reflect the name of the user to whom the API privileges should be granted.
- b. Adjust the name of the container's API schema C#DI to suit your needs.
- c. If the target user should also be able to grant another user the container administration privileges, replace the procedure C#DI.GRANT_CONTAINER_API_PRIVILEGES in the SQL statement above with the procedure C#DI.GRANT_CONTAINER_API_PRIVILEGES_WITH_GRANT_OPTION.

Note

The procedure C#DI.GRANT_CONTAINER_API_PRIVILEGES_WITH_GRANT_OPTION should only be used in special scenarios, for example, when building a custom SQL API by wrapping the HDI SQL API in SQLScript procedures.

4. Execute the SQL code.
Confirm that the SQL code completes successfully and displays the HDI return code 0.
5. (Optional) Confirm that the NEW_CONTAINER_ADMIN user is now able to call HDI container API procedures in the containers API schema (for example, C#DI in container "C").

Related Information

[Revoke HDI Container Administrator Privileges from a User \[page 1511\]](#)

[Maintaining HDI Containers \[page 1508\]](#)

11.7.2 Revoke HDI Container Administrator Privileges from a User

Disable administrator access to an HDI container.

Context

Each container can have its own set of administrators. Administrative privileges for a container must be explicitly granted or revoked by an HDI container group administrator, or an HDI container administrator with the necessary privileges.

Note

The predefined table `_SYS_DI.T_DEFAULT_CONTAINER_ADMIN_PRIVILEGES` used in this task contains the largest possible set of privileges that can be granted for a user of this type. You can reduce the set of privileges granted by explicitly specifying the desired set of privileges and not using this default table.

Procedure

1. In an SQL console, connect to the database with an administrator of the HDI container “C” (or the name of the container whose administrator privileges you want to change).
2. Open the SQL editor for this database.
3. Revoke container-administrator privileges from a user.

Insert the following SQL code into the SQL editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
_SYS_DI.TT_API_PRIVILEGES;
INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME)
SELECT 'NEW_CONTAINER_ADMIN', PRIVILEGE_NAME, OBJECT_NAME FROM
_SYS_DI.T_DEFAULT_CONTAINER_ADMIN_PRIVILEGES WHERE NOT (PRIVILEGE_NAME =
'SELECT' AND OBJECT_NAME LIKE '_SYS_DI.T%');
CALL _SYS_DI#G.REVOKE_CONTAINER_API_PRIVILEGES('C', #PRIVILEGES,
_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;
```

- a. In the `INSERT` statement in line 2, adjust the name of the user “`NEW_CONTAINER_ADMIN`” to reflect the name of the user from whom the API privileges should be revoked.
 - b. In the `CALL` statement in line 3, adjust the name of the container's API schema `C#DI` to suit your needs.
4. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

5. (Optional) Confirm that the `NEW_CONTAINER_ADMIN` user is no longer able to call HDI container API procedures in the containers API schema (for example, `C#DI` in container “C”).

Related Information

[Maintaining HDI Containers \[page 1508\]](#)

[Grant HDI Container Administrator Privileges to a User \[page 1509\]](#)

11.7.3 Grant Access to the HDI Container Content-Development API

Enable access to the HDI-container, content-development, application programming interface (API).

Context

The content-development API in a container's #DI schema (for example, schema C#DI in container "C") is intended for use by developers of HDI database artifacts. Developers use the container content-development API to write design-time artifacts to the container; the design-time artifacts are then used to generate database objects in the container's schema (for example, schema C in container "C"). A container administrator must first grant the necessary privileges to the developer before the API can be used, as described below.

Note

The predefined table `_SYS_DI.T_DEFAULT_CONTAINER_USER_PRIVILEGES` used in this task contains the largest possible set of privileges that can be granted for a user of this type. You can reduce the scope of the privileges granted by not using this default table and explicitly specifying the desired set of privileges instead.

Procedure

1. In an SQL console, connect to the database with an administrator of the target HDI container (for example, "C").
2. Open the SQL editor for this database.
3. Grant access to the HDI container content-development API.

Insert the following SQL code into the editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
_SYS_DI.TT_API_PRIVILEGES;
INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME)
SELECT 'NEW_CONTAINER_CONTENT_DEVELOPER', PRIVILEGE_NAME, OBJECT_NAME FROM
_SYS_DI.T_DEFAULT_CONTAINER_USER_PRIVILEGES;
CALL C#DI.GRANT_CONTAINER_API_PRIVILEGES(#PRIVILEGES,
_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

```
DROP TABLE #PRIVILEGES;
```

- a. Adjust the name of the new content-development user “NEW_CONTAINER_CONTENT_DEVELOPER” in the INSERT statement in line 2 to reflect the name of the user to whom the API privileges should be granted.
 - b. Adjust the name of the container's API schema C#DI in the CALL statement to suit your needs.
4. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

5. (Optional) Confirm that the NEW_CONTAINER_CONTENT_DEVELOPER user is now able to call the HDI container content-development API in the container's API schema (for example, C#DI in container “C”).

Related Information

[Maintaining HDI Containers \[page 1508\]](#)

[Revoke Access to the HDI Container Content-Development API \[page 1513\]](#)

11.7.4 Revoke Access to the HDI Container Content-Development API

Disable access to the HDI-container, content-development, application programming interface (API).

Context

The container content-development API in a container's schema (for example, schema C#DI in container “C”) is intended for use by developers of HDI database artifacts. Developers use the container content-development API to write design-time artifacts to the container; the design-time artifacts are then used to generate database objects in the container's schema (for example, schema C#DI in container “C”). A container administrator can revoke the privileges that enable access to the container content-development API at any time, as described below.

i Note

The predefined table `_SYS_DI.T_DEFAULT_CONTAINER_USER_PRIVILEGES` used in this task contains the largest possible set of privileges that can be granted for a user of this type. You can reduce the scope of the privileges granted by not using this default table and explicitly specifying the desired set of privileges instead.

Procedure

1. In an SQL console, connect to the database as an administrator of the HDI container whose development API you want to enable (for example, "C").
2. Open the SQL editor for this database.
3. Revoke access to the HDI container content-development API.

Insert the following SQL code into the editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
SYS_DI.TT_API_PRIVILEGES;
INSERT INTO #PRIVILEGES (PRINCIPAL_NAME, PRIVILEGE_NAME, OBJECT_NAME)
SELECT 'NEW_CONTAINER_CONTENT_DEVELOPER', PRIVILEGE_NAME, OBJECT_NAME FROM
SYS_DI.T_DEFAULT_CONTAINER_USER_PRIVILEGES WHERE NOT (PRIVILEGE_NAME =
'SELECT' AND OBJECT_NAME LIKE 'SYS_DI.T%');
CALL C#DI.REVOKE_CONTAINER_API_PRIVILEGES(#PRIVILEGES,
SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;
```

- a. In the `INSERT` statement in line 2, adjust the name of the new content-development user "NEW_CONTAINER_CONTENT_DEVELOPER" to reflect the name of the user from whom the API privileges should be revoked.
 - b. In the `CALL` statement in line 3, adjust the name of the container's API schema `C#DI` to suit your needs.
4. Execute the SQL code.
Confirm that the SQL code completes successfully and displays the HDI return code 0.
 5. (Optional) Confirm that the `NEW_CONTAINER_CONTENT_DEVELOPER` user is no longer able to call the HDI container content-development API in the container's API schema (for example, `C#DI` in container "C").

Related Information

[Maintaining HDI Containers \[page 1508\]](#)

[Grant Access to the HDI Container Content-Development API \[page 1512\]](#)

11.7.5 Grant Access to an HDI Container's Schema

Enable access to the schema of an HDI container of individual objects in the target schema.

Context

Users that would like to consume objects deployed to an HDI container need to be granted the appropriate privileges. The privileges can be granted to specific objects in the schema (for example, c) by use of a role that has been deployed to the target container, for example, "C", or by granting privileges for the entire schema.

To grant access privileges for the entire container schema where the database objects are located to a database object consumer `NEW_CONTAINER_CONSUMER`, perform the following steps:

Procedure

1. In an SQL console, connect to the database with an administrator of the target HDI container (for example, "C").
2. Open the SQL editor for this database.
3. Grant access to the HDI container's schema.

Insert the following SQL code into the editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
SYS_DI.TT_SCHEMA_PRIVILEGES;
INSERT INTO #PRIVILEGES ( PRIVILEGE_NAME, PRINCIPAL_SCHEMA_NAME,
PRINCIPAL_NAME ) VALUES ( 'SELECT', '', 'NEW_CONTAINER_CONSUMER' );
CALL C#DI.GRANT_CONTAINER_SCHEMA_PRIVILEGES ( #PRIVILEGES,
SYS_DI.T_NO_PARAMETERS, ?, ?, ? );
DROP TABLE #PRIVILEGES;
```

- a. Adjust the name of the consumer user "NEW_CONTAINER_CONSUMER" in the INSERT statement in the INSERT statement in line 2 to reflect the name of the user who requires access to the container.
 - b. Adjust the name of the container's API schema (C#DI) in the CALL statement to suit your needs.
4. Execute the SQL code.
Confirm that the SQL code completes successfully and displays the HDI return code 0.
 5. (Optional) Confirm that the NEW_CONTAINER_CONSUMER can now access the database objects in the container's schema.

Related Information

[Maintaining HDI Containers \[page 1508\]](#)

11.7.6 Revoke Access to an HDI Container's Schema

Revoke access to the schema of an HDI container or specific objects in the target schema.

Context

Users that need to consume objects deployed in a container (for example, "C") must be granted the appropriate privileges to access the target container's schema (for example, C#DI). The privileges for the entire schema can be revoked with the `REVOKE_CONTAINER_SCHEMA_PRIVILEGES` API.

To revoke privileges for the entire container schema where the database objects are located from a database object consumer `NEW_CONTAINER_CONSUMER` perform the following steps:

Procedure

1. In an SQL console, connect to the database with an administrator of the target HDI container (for example, "C").
2. Open the SQL editor for this database.
3. Revoke access to the HDI container's schema.

Insert the following SQL code into the editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PRIVILEGES LIKE
SYS_DI.TT_SCHEMA_PRIVILEGES;
INSERT INTO #PRIVILEGES ( PRIVILEGE_NAME, PRINCIPAL_SCHEMA_NAME,
PRINCIPAL_NAME ) VALUES ( 'SELECT', '', 'NEW_CONTAINER_CONSUMER' );
CALL C#DI.REVOKE_CONTAINER_SCHEMA_PRIVILEGES( #PRIVILEGES,
SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #PRIVILEGES;
```

- a. In the `INSERT` statement in line 2, adjust the name of the consumer user "NEW_CONTAINER_CONSUMER" to reflect the name of the user whose container-access privileges must be revoked.
 - b. In the `CALL` statement, adjust the name of the container's API schema (C#DI) to suit your needs.
4. Execute the SQL code.
Confirm that the SQL code completes successfully and displays the HDI return code 0.
 5. (Optional) Confirm that the `NEW_CONTAINER_CONSUMER` can no longer access the database objects in the container's schema.

Related Information

[Maintaining HDI Containers \[page 1508\]](#)

[Grant Access to an HDI Container's Schema \[page 1515\]](#)

[HDI Container Schemas \[page 1468\]](#)

11.7.7 Grant a User a Role from the Container's Schema

Enable access to schema objects by means of a role.

Context

Users that would like to consume objects deployed to a container (for example, "C") need to be granted the appropriate privileges. The access privileges can be granted to specific objects in the schema (for example, C#DI) by use of a role that has been deployed to the container, or by granting privileges for the entire schema.

To grant privileges for specific objects in the container to a user `NEW_CONTAINER_CONSUMER` by use of a role in the target container, perform the following steps:

Procedure

1. In an SQL console, connect to the database with an administrator of the target HDI container (for example, "C").
2. Open the SQL editor for this database.
3. Grant the role to the target user.

Insert the following SQL code into the editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #ROLES LIKE _SYS_DI.TT_SCHEMA_ROLES;  
INSERT INTO #ROLES ( ROLE_NAME, PRINCIPAL_SCHEMA_NAME, PRINCIPAL_NAME )  
VALUES ( 'myrole', '', 'NEW_CONTAINER_CONSUMER' );  
CALL C#DI.GRANT_CONTAINER_SCHEMA_ROLES(#ROLES,  
_SYS_DI.T NO PARAMETERS, ?, ?, ?);  
DROP TABLE #ROLES;
```

- a. Adjust the name of the consumer user `NEW_CONTAINER_CONSUMER` in the `INSERT` statement in line 2
 - b. Adjust the name of the role "myrole" in the `INSERT` statement in line 2.
 - c. Adjust the schema name of the container C's API schema `C#DI` in the `CALL` statement in line 3.
4. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

5. (Optional) Confirm that the `NEW_CONTAINER_CONSUMER` can now access the database objects specified by the role in the container's schema.

Related Information

[Maintaining HDI Containers \[page 1508\]](#)

[Revoke a Role from the Container's Schema \[page 1518\]](#)

11.7.8 Revoke a Role from the Container's Schema

Disable access to schema objects by means of a role.

Context

Users that would like to consume objects deployed to a container (for example, "C") need to be granted the appropriate privileges. The privileges for specific objects in the target schema (for example, `C#DI`), granted by a role in the schema, can be revoked with the `REVOKE_CONTAINER_SCHEMA_ROLES` API:

To revoke a role and disable access privileges for specific objects in the container from a user, perform the following steps:

Procedure

1. In an SQL console, connect to the database with an administrator of the target HDI container (for example, "C").
2. Open the SQL editor for this database.
3. Grant the role to the target user.

Insert the following SQL code into the editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #ROLES LIKE _SYS_DI.TT_SCHEMA_ROLES;
INSERT INTO #ROLES ( ROLE_NAME, PRINCIPAL_SCHEMA_NAME, PRINCIPAL_NAME )
VALUES ( 'myrole', '', 'NEW_CONTAINER_CONSUMER' );
CALL C#DI.REVOKE_CONTAINER_SCHEMA_ROLES(#ROLES,
SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #ROLES;
```

- a. In the `INSERT` statement in line 2, adjust the name of the consumer user `NEW_CONTAINER_CONSUMER` and the role "myrole" to suit your needs.

- b. In the `CALL` statement in line 3, adjust the schema name of the container "C"'s API schema `C#DI`.
4. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

5. (Optional) Confirm that the `NEW_CONTAINER_CONSUMER` can no longer access the database objects specified by the role in the container's schema.

Related Information

[Maintaining HDI Containers \[page 1508\]](#)

[Grant a User a Role from the Container's Schema \[page 1517\]](#)

11.7.9 List All Currently Configured Build Plug-in Libraries Available to a Container

Display a list of all the build plug-in libraries available for use in a container.

Context

Administrators of an HDI container (for example, container "C") can view all currently configured HDI build plug-in libraries and their versions for any container where they have the container-administration permissions, as described in the following procedure:

Procedure

1. In an SQL console, connect to the database with an administrator of the target HDI container, for example, "C".
2. Open the SQL editor for this database.
3. Display a list of all currently available plug-ins in the specified container.

Insert the following SQL statement into the SQL editor:

Sample Code

```
CALL C#DI.LIST_CONFIGURED_LIBRARIES (_SYS_DI.T_NO_PARAMETERS, ?, ?, ?, ?);
```

- a. Replace the name of the container "C" in the `CALL` statement with the name of your container.
4. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

5. A result set with the requested plug-in information is returned.

The following example shows an excerpt from the result set returned:

Output Code

```
LIBRARY_NAME;LIBRARY_VERSION;PLUGIN_ID;PLUGIN_VERSION
com.sap.hana.di.afllangprocedure;0.0;com.sap.hana.di.afllangprocedure;
2.0.10.0
com.sap.hana.di.analyticprivilege;0.0;com.sap.hana.di.analyticprivilege;
2.0.10.0
com.sap.hana.di.calculationview;0.0;com.sap.hana.di.calculationview;
2.0.10.0
com.sap.hana.di.cds;0.3;com.sap.hana.di.cds;2.0.10.0
com.sap.hana.di.constraint;0.0;com.sap.hana.di.constraint;2.0.10.0
com.sap.hana.di.copyonly;0.0;com.sap.hana.di.copyonly;2.0.10.0
com.sap.hana.di.dropcreatetable;0.0;com.sap.hana.di.dropcreatetable;
2.0.10.0
com.sap.hana.di.flowgraph;0.0;com.sap.hana.di.flowgraph;2.0.10.0
```

Related Information

[Configure the Default Build Plug-in Libraries Available to a Container \[page 1520\]](#)

[Configure a Custom Set of Build Plug-in Libraries Available to a Container \[page 1521\]](#)

[Maintaining HDI Containers \[page 1508\]](#)

11.7.10 Configure the Default Build Plug-in Libraries Available to a Container

Maintain the set of plug-in libraries available by default in an HDI container.

Context

HDI container administrators can configure a default set of commonly used HDI plug-in libraries for any HDI container for which they have responsibility.

Procedure

1. In an SQL console, connect to the database as an administrator of the target HDI container, for example, "C".
2. Open the SQL editor for this database.
3. Set the list of plug-in libraries available by default in the specified target container.

Insert the following SQL statement into the SQL editor:

Sample Code

```
CALL C#DI.CONFIGURE_LIBRARIES(_SYS_DI.T_DEFAULT_LIBRARIES,  
_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

- a. Replace the name of the container "C" in the `CALL` statement with the name of your container.
4. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.
5. **(Optional)** Confirm that the default libraries have been configured successfully by listing all currently configured build plug-in libraries available to the container.

Related Information

[List All Currently Configured Build Plug-in Libraries Available to a Container \[page 1519\]](#)

[Configure a Custom Set of Build Plug-in Libraries Available to a Container \[page 1521\]](#)

[Maintaining HDI Containers \[page 1508\]](#)

11.7.11 Configure a Custom Set of Build Plug-in Libraries Available to a Container

Maintain a custom set of plug-in libraries available in an HDI container.

Context

HDI container administrators can configure a custom set of commonly used HDI plug-in libraries for any container for which they are responsible. It is possible to define the full set of libraries, or incrementally add or remove certain libraries to the currently configured set.

Procedure

1. In an SQL console, connect to the database as an administrator of the target HDI container, for example, "C".
2. Open the SQL editor for this database.
3. Define a custom list of plug-in libraries available in the specified target container.

Insert the following SQL statement into the SQL editor:

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #LIBRARY_CONFIGURATION LIKE
_SYS_DI.TT_LIBRARY_CONFIGURATION;
INSERT INTO #LIBRARY_CONFIGURATION ( ACTION, LIBRARY_NAME ) VALUES
( 'ADD', 'com.sap.hana.di.calculationview' );
INSERT INTO #LIBRARY_CONFIGURATION ( ACTION, LIBRARY_NAME ) VALUES
( 'ADD', 'com.sap.hana.di.cds' );
INSERT INTO #LIBRARY_CONFIGURATION ( ACTION, LIBRARY_NAME ) VALUES
( 'ADD', 'com.sap.hana.di.synonym' );
INSERT INTO #LIBRARY_CONFIGURATION ( ACTION, LIBRARY_NAME ) VALUES
( 'REMOVE', 'com.sap.hana.di.view' );
CALL C#DI.CONFIGURE_LIBRARIES(#LIBRARY_CONFIGURATION,
_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #LIBRARY_CONFIGURATION;
```

- a. Replace the name of the container “C” in the `CALL` statement with the name of your container.
 - b. Adjust the set of libraries to be added or removed in the `INSERT` statements as needed. This example adds three libraries and removes one.
4. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

5. **(Optional)** Confirm that the custom set of libraries has been configured successfully by listing all currently configured build plug-in libraries available in the container.
6. **(Optional)** Customize the configuration of the build plug-in library if necessary.

The behavior of the build plug-ins library configuration can be modified by supplying custom parameters in the call to `CONFIGURE_LIBRARIES` instead of `_SYS_DI.T_NO_PARAMETERS`.

For example, to remove a build plug-in library and **undeploy** all files ('undeploy', 'true') corresponding to that library, use the following SQL code:

Note

The default value for the `undeploy` parameter is 'false'

Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #PARAMETERS LIKE _SYS_DI.TT_PARAMETERS;
INSERT INTO #PARAMETERS (KEY, VALUE) VALUES ('undeploy', 'true');
CREATE LOCAL TEMPORARY COLUMN TABLE #LIBRARY_CONFIGURATION LIKE
_SYS_DI.TT_LIBRARY_CONFIGURATION;
INSERT INTO #LIBRARY_CONFIGURATION ( ACTION, LIBRARY_NAME ) VALUES
( 'REMOVE', 'com.sap.hana.di.view' );
CALL C#DI.CONFIGURE_LIBRARIES(#LIBRARY_CONFIGURATION,
#PARAMETERS, ?, ?, ?);
DROP TABLE #LIBRARY_CONFIGURATION;
DROP TABLE #PARAMETERS;
```

Related Information

[List All Currently Configured Build Plug-in Libraries Available to a Container \[page 1519\]](#)

11.7.12 Configure Container Parameters

Maintain the parameters used to configure an HDI container.

Context

The HDI container administrator can use configuration parameters to maintain some of the more general aspects of an HDI container. Container-specific configuration parameters are used to control the behavior of a single container, for example, they can be used to specify the time a container operation waits for a locking conflict to clear, or the maximum number of parallel jobs to be spawned during a make operation.

→ Tip

For more information about HDI configuration parameters, see *Related Information*.

To configure container-specific parameters, perform the following steps:

Procedure

1. In an SQL console, connect to the database with an administrator of the target HDI container (for example, "C").
2. Open the SQL editor for this database.
3. Define and set the parameters required for the configuration of the target HDI container.

≡ Sample Code

```
CREATE LOCAL TEMPORARY COLUMN TABLE #CONFIG_PARAMETERS LIKE
_SYS_DI.TT_PARAMETERS;
INSERT INTO #CONFIG_PARAMETERS (KEY, VALUE) VALUES
('make.max_parallel_jobs', '8');
CALL C#DI.CONFIGURE_CONTAINER_PARAMETERS(#CONFIG_PARAMETERS,
_SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
DROP TABLE #CONFIG_PARAMETERS;
```

- a. Add more parameters and values as required.

→ Tip

A new `INSERT` statement is required for each new parameter.

```
INSERT INTO #CONFIG_PARAMETERS (KEY, VALUE) VALUES
('make.max_parallel_jobs', '8');
```

```
INSERT INTO #CONFIG_PARAMETERS (KEY, VALUE) VALUES
('connection.transaction_lock_wait_timeout', '1,000');
INSERT INTO #CONFIG_PARAMETERS (KEY, VALUE) VALUES
('messages.days_to_keep', '30');
...
```

- b. Adjust the schema name of the container ("C") in the `CALL` statement.
4. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

Related Information

[SAP HANA DI Parameters \[page 1469\]](#)

[SAP HANA DI Configuration Parameters \[page 1484\]](#)

11.7.13 Cancel a Running Make Operation in a Container

Context

An HDI container administrator can cancel a running make in a container C, for example, if it takes longer than expected. This is mainly intended for canceling asynchronous makes (C#DI.MAKE_ASYNC) but also works for synchronous ones (C#DI.MAKE). When an asynchronous make is started, it returns a request ID identifying that make process. This request ID can then be used in a call to the C#DI.CANCEL procedure for canceling this process.

Procedure

1. In an SQL console, connect to the database as the administrator of the target HDI container (for example, "C").
2. Open the SQL editor for this database.
3. Cancel the make operation.

Enter the following SQL code:

Sample Code

```
CALL C#DI.CANCEL(12345, _SYS_DI.T_NO_PARAMETERS, ?, ?, ?);
```

- a. Adjust the name of the container ("C") in the `CALL` statement to reflect your target container.

- b. Replace “12345” in the `CALL` statement with the request ID of the make process you want to cancel.

→ Tip

The procedure `C#DI.MAKE` returns one or more result sets, which contain the request ID and the return code.

4. Execute the SQL code.

Confirm that the SQL code completes successfully and displays the HDI return code 0.

Related Information

[Maintaining HDI Containers \[page 1508\]](#)

12 Application Run-Time Services

Maintain the SAP HANA XS run-time environment for XS classic and XS advanced applications.

The SAP HANA administration cockpit provides the tools you need to maintain and manage the various components of the SAP HANA XS run-time environment. Whether you are providing administration and support services for applications running in the XS classic run time or you need to set up and maintain an XS advanced run time in SAP HANA, the administration cockpit provides a selection of tools to help you perform your tasks quickly and easily.

- SAP HANA XS classic model
Maintain and manage the various components of the SAP HANA XS classic Model (XS classic) run-time environment
- SAP HANA XS advanced model
Maintain and manage the various components of the SAP HANA XS Advanced Model (XS advanced) run-time environment

Related Information

[Maintaining the SAP HANA XS Classic Model Run Time \[page 1526\]](#)

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1647\]](#)

12.1 Maintaining the SAP HANA XS Classic Model Run Time

Maintain the SAP HANA XS classic model run-time environment.

A number of administration tools are available to enable you to maintain and manage the various components of the SAP HANA XS classic model (XS classic) run-time environment. In the SAP HANA administration cockpit, the *XS Administration* tile catalog contains the *Administration and Monitoring* tile, which contains the following tools:

i Note

In the SAP HANA cockpit, tiles and tile catalogs are only visible to users who have been assigned the privileges granted by role `sap.hana.uis.db::SITE_DESIGNER`. In addition, some of the tools listed below are only available to users to whom the suitable role has been assigned. For example, a role based on the role template `sap.hana.xs.admin.roles::RuntimeConfAdministrator` includes the authorization required for unrestricted access to all the tools used to manage the configuration settings for SAP HANA XS application security and the related user-authentication providers; a role based on the role template `sap.hana.xs.admin.roles::SAMLAdministrator` enables unrestricted access only to the *SAML Identity Providers Configuration* tools.

- [XS Artifact Administration](#)
Monitor the system usage of the applications running in the XS Advanced Model run-time
- [SAML Service Provider](#)
Configure an SAP HANA system to act as an SAML service provider for SSO authentication.
- [SAML Identity Provider](#)
Configure an SAML identity provider for use by the SAML service provider to authenticate the users signing in by means of SSO.
- [SMTP Configuration](#)
Maintain and manage details of the SMTP server that is available for use by all applications running on an SAP HANA XS classic model server.
- [Trust Manager](#)
Configure SAML Identity providers (IDP) for SAP HANA XS classic model applications that use SAML assertions as the log-on authentication method.
- [XS Job Dashboard](#)
Create, schedule, and manage long running operations jobs in the SAP HANA XS classic model run-time environment.

Related Information

[SAP HANA XS Classic Administration Tools \[page 1527\]](#)

[SAP HANA XS Classic Administration Roles \[page 1529\]](#)

[SAP HANA XS Classic Configuration Parameters \[page 1532\]](#)

12.1.1 SAP HANA XS Classic Administration Tools

SAP HANA XS includes a Web-based tool that enables you to maintain important parts of the application-development environment, for example, security and authentication methods.

The *SAP HANA XS Administration Tool* is a Web-based tool that enables you to configure and maintain the basic administration-related elements of the application-development process and environment. The features included in the Web-based *SAP HANA XS Administration Tool* cover the following areas:

i Note

The availability of screens, tabs, and UI controls (for example, *Add*, *Edit*, or *Save* buttons) is based on the privileges granted in the assigned user roles. For example, a user who has a role based on the role template `sap.hana.xs.admin.roles::HTTPDestViewer` can view HTTP destinations; a user assigned a role based on the role template `sap.hana.xs.admin.roles::SQLCCAdministrator` can not only view but also **edit** SQL connection configurations.

Administration Tools for SAP HANA XS Applications

Tool Name	Description	Scope
<i>XS Artifact Administration</i>	Maintain runtime configurations for individual applications or a complete application hierarchy. The configuration defined for an application is inherited by any application further down the application package hierarchy.	<ul style="list-style-type: none"> • Application security (public/private) • User-authentication methods (basic, form-based, logon tickets, X509, SAML) • CORS setup for cross-origin resource sharing • Custom headers: enable support for X-Frame-Options HTTP header • HTTP destinations • SQL connection configurations (for SQL connections for users other than the user specified in the HTTP request).
<i>SAML Service Provider</i>	Configure an SAP HANA system to act as an SAML service provider for SSO authentication.	<ul style="list-style-type: none"> • Management of SAML service-providers, including URLs and meta-data management
<i>SAML Identity Provider</i>	Configure an SAML identity provider for use by the SAML service provider to authenticate the users signing in by means of SSO.	<ul style="list-style-type: none"> • Management of SAML identity-providers, including IDP metadata, certificates, and destinations
<i>SMTP Configuration</i>	Define the details of the SMTP server that is available for use by all applications running on an SAP HANA XS server.	<ul style="list-style-type: none"> • SMTP host settings • Authentication type • Transport security
<i>Trust Manager</i>	Maintain the certificates used to establish trust relationships between servers used by SAP HANA XS applications.	<ul style="list-style-type: none"> • Trust store configuration and management • Certificate management
<i>XS Job Dashboard</i>	Monitor and maintain SAP HANA XS job schedules defined using the XS job syntax	<ul style="list-style-type: none"> • Enable the job scheduler • Monitor job-schedule status • Display and maintain schedule's runtime configuration • Add schedules to (or delete from) an XS job

Additional Tools in SAP HANA XS Classic

The following table lists some tools that are not strictly part of the SAP HANA XS Administration tool set. The tools are included here primarily for the sake of convenience but also because the tools are installed with the delivery unit which contains the XS Administration tools.

Translation Text Details

Tool Name	Description	Scope
Online Translation Tool	Maintain translations, for example, for UI text elements	<ul style="list-style-type: none">• Add, modify, delete translation texts• Export translation text from SAP HANA to an XML-based xliif-format file• Import translation text into SAP HANA.
User Self Service Tools	A set of tools that enable you to maintain user self-service requests and administrate the self-service tools themselves.	<ul style="list-style-type: none">• Activate the user self-service tools• Maintain user self-service requests• Maintain user black/white lists• Maintain user self-service e-mail templates

Related Information

[Maintaining Application Runtime Configurations \[page 1538\]](#)

[Maintaining SAML Providers \[page 1559\]](#)

[Managing Trust Relationships \[page 1552\]](#)

[Maintaining SMTP Server Configurations \[page 1568\]](#)

[Scheduling XS Jobs \[page 1618\]](#)

[Maintaining User Self Service Tools \[page 1593\]](#)

[Maintaining Translation Text Strings \[page 1632\]](#)

12.1.2 SAP HANA XS Classic Administration Roles

SAP HANA uses roles to control access to the Web-based tool that enable you to maintain important parts of the application-development environment, for example, security and authentication methods.

When using the Web-based tools provided by SAP HANA XS, the availability of features, screens, tabs, and UI controls (for example, *Add*, *Edit*, or *Save*, or *Delete* buttons) is based on privileges. For the sake of convenience, the specific privileges required to use the features provided with a particular tool have been collected into a selection of predefined roles, which you can use as templates to create your own roles and assign to the user who wants to use a tool. For example, a user assigned a role based on `sap.hana.xs.admin.roles::HTTPDestViewer` can display HTTP destinations but not change them in any way; a user assigned a role based on `sap.hana.xs.admin.roles::SQLCCAdministrator` can view SQL connection configurations and modify them, too.

→ Recommendation

As repository roles delivered with SAP HANA can change when a new version of the package is deployed, either do not use them directly but instead as a template for creating your own roles, or have a regular review process in place to verify that they still contain only privileges that are in line with your organization's

security policy. Furthermore, if repository package privileges are granted by a role, we recommend that these privileges be restricted to your organization's packages rather than the complete repository. To do this, for each package privilege (`REPO.*`) that occurs in a role template and is granted on `.REPO_PACKAGE_ROOT`, check whether the privilege can and should be granted to a single package or a small number of specific packages rather than the full repository.

SAP HANA XS Administration Tools Roles

SAP HANA XS Role	Description
HTTPDestAdministrator	Full access to the details of HTTP destination configurations (display and edit)
HTTPDestViewer	Read-only access to HTTP destination configurations, which are used to specify connection details for outbound connections, for example, using the server-side JavaScript Connectivity API that is included with SAP HANA XS.
RuntimeConfAdministrator	Full access to the configuration settings for SAP HANA XS application security and the related user-authentication providers.
RuntimeConfViewer	Read-only access to the configuration settings for SAP HANA XS application security and the related user-authentication providers, for example, SAML or X509.
JobAdministrator	Full access to the configuration settings for SAP HANA XS job schedules (defined in <code>.xsjob</code> files); you can specify start/stop times, the user account to run the job, and the language locale.
JobViewer	Read-only access to the configuration settings for SAP HANA XS job schedules (defined in <code>.xsjob</code> files).
JobScheduleAdministrator	Full access to the <i>XS Job Dashboard</i> tool, which you can use to add and delete XS job schedules, maintain individual schedules, and enable the scheduling feature.
oAuthAdmin	Required when setting the client secret during administration of the OAuth client configuration (<code>.xsoauthclientconfig</code>) artifact.
SAMLAdministrator	Full access to the details of SAML configurations, including both the service provider and the identity providers. You can add new entries and make changes to existing service or identity providers and parse the resulting metadata.
SAMLViewer	Read-only access to SAML configurations, which are used to provide details of SAML service providers and identity providers.
SMTPDestAdministrator	Full access to the details of SMTP destination configurations, which are used to define details of the SMTP relay server that SAP HANA XS applications use to send e-mails. The administrator role enables you to add new entries and make changes to an existing configuration, for example, the host name and port number, logon credentials and authentication type, and any transport security settings.
SMTPDestViewer	Read-only access to SMTP destination configurations, which are used to define details of the SMTP relay server that SAP HANA XS applications can use to send e-mails.
SQLCCAdministrator	Full access to the details of SQL connection configurations (SQLCC).
SQLCCViewer	Read-only access to SQL connection configurations (SQLCC), which are used to enable the execution of SQL statements from inside your server-side JavaScript application with credentials that are different to the credentials of the requesting user.
TrustStoreAdministrator	Full access to the SAP HANA XS <i>Trust Manager</i> tool, which the administrator uses to maintain secure outbound communication, for example, the SSL/TLS certificates required by SAP HANA XS applications that connect to an ABAP system.

SAP HANA XS Role	Description
TrustStoreViewer	Read-only access to the trust store, which contains the server's root certificate or the certificate of the certification authority that signed the server's certificate.

Additional SAP HANA XS Roles

The following table lists roles for tools that are not strictly part of the SAP HANA XS Administration toolset. The roles are included here for the sake of convenience and because the roles, and the tools to which they correspond, are (with the exception of the *WebDispatcherAdmin/Viewer*) installed with the delivery unit which contains the XS Administration tools.

Additional Roles for SAP HANA XS Administration Tools

SAP HANA XS Role	Description
translator	The role <i>sap.hana.xs.translationTool.roles::translator</i> enables an SAP HANA user to maintain translation text strings with the SAP HANA Online Translation Tool.
USSAdministrator	The role <i>sap.hana.xs.selfService.admin.roles::USSAdministrator</i> is assigned to the user responsible for administrating the requests sent by users using self-service tools. For example, it enables the activation of users who request a new user account in the SAP HANA database and allows the user-self-service administrator to manage self-service-specific blacklists for users, e-mail addresses, domains, and IP addresses.
USSExecutor	The role <i>sap.hana.xs.selfService.user.roles::USSExecutor</i> is assigned to the technical user that is used to respond to and execute user-self-service requests, for example, to create a new account or request a new password.
WebDispatcherAdmin	The role <i>sap.hana.xs.wdisp.admin::WebDispatcherAdmin</i> enables full access to the SAP HANA <i>Web Dispatcher Administration</i> tool, which the administrator uses to maintain secure inbound communication, for example, to enable SSL/TLS connections between an ABAP system and an SAP HANA XS application.
WebDispatcherMonitor	The role <i>sap.hana.xs.wdisp.admin::WebDispatcherMonitor</i> enables read-only access to the information displayed in the SAP HANA <i>Web Dispatcher Administration</i> tool.
WebDispatcherHTTPTracingViewer	Read-only access to the HTTP setting of SAP HANA XS applications running on the selected SAP HANA instance. This role extends the <i>JobViewer</i> role to enable the user to view details of the <i>xsjob</i> configuration (<i>httptracing.xsjob</i>) that starts and stops the HTTP tracing tasks.
WebDispatcherHTTPTracingAdministrator	Full access required to maintain HTTP tracing on the SAP Web Dispatcher for SAP HANA XS applications. This role extends the <i>JobAdministrator</i> role to enable the user to maintain the XS job file (<i>httptracing.xsjob</i>) used to configure and enable HTTP tracing for XS applications on the SAP Web Dispatcher.

12.1.3 SAP HANA XS Classic Classic Configuration Parameters

An overview of the parameters that the administrator can set to configure how the various components of the XS engine work.

The `xsengine.ini` section of the SAP HANA configuration screen is split into a number of subsections, each of which reflects one of the individual components of the SAP HANA XS engine. Each section contains one or more parameters whose values you can change, where appropriate, to suit the requirements of your system landscape. To display the configuration details of the XS engine in SAP HANA studio, double-click a system in the *Systems* view, choose the *Configuration* tab, and expand the *xsengine.ini* element.

Note

For security reasons, all parameters in the `communication` section of **all** `.ini` configuration files are blacklisted by default; properties included in a blacklist can only be changed by a system administrator. For more information, see *Default Blacklisted System Properties* in the *SAP HANA Administration Guide*.

XS Engine Configuration Parameters (xsengine.ini)

Configuration Section	Description
application_container [page 1533]	Application-related configuration settings, for example, the list of applications that are trusted by the XS engine or the libraries that can be loaded from an <code>xsconfunc</code> call.
authentication [page 1533]	Options for application-related authentication settings, for example, the location of trust stores.
communication [page 1533]	Options for application-related connection requests and configuration, for example, time-outs, port numbers, and maximum number of data end points allowed by the XS engine
customer_usage [page 1534]	Options for customer-specific usage scenarios in SAP HANA application services, for example, to enable HTTP tracing of XS applications on the SAP Web Dispatcher.
debugger [page 1534]	Settings for the debugging tools, for example, for XS JavaScript.
httpservlet [page 1535]	Options for the SAP HANA XS Web server, for example, port numbers, and maximum number of sessions and threads allowed
odata [page 1536]	Configuration settings for OData requests
scheduler [page 1536]	Configuration options for the XS job scheduler, which is used to run an XS Javascript or SQLScript as a task in the background at regular intervals

Note

Some configuration parameters for the SAP HANA XS engine require additional parameters to be set for other SAP HANA components, for example, the SAP Web Dispatcher.

SAP Web Dispatcher Configuration Parameters (webdispatcher.ini)

Configuration Section	Description
webdispatcher.ini/profile [page 1536]	Configuration options for the SAP Web Dispatcher, for example: HTTP tracing of SAP HANA XS applications, logs, allowed connections

application_container

Use the `application_container` section of the `xsengine.ini` file to set configuration options for the application container component of the SAP HANA XS engine, which includes not only the XS application container, but also containers for C++ and JavaScript applications. In this section of the `xsengine.ini` file, you can modify the list of applications that are trusted by the XS engine or the libraries that can be loaded from an `xscfunc` call.

Parameter	Description	Example Value	Default Value
<code>application_list</code>	Comma-separated list of libraries that can be loaded from an <code>xscfunc</code> call	<code>libxsdxs,</code> <code>InformationAccess</code>	<code>libxsdxs,</code> <code>InformationAccess,</code> <code>libtrustmanager,</code> <code>libxsauthenticator,</code> <code>libxsbase</code>

authentication

Use the `authentication` section of the `xsengine.ini` file to set configuration options for application-related authentication settings, for example, the system ID and hostname of the server providing SAP logon certificates for single sign-on (SSO) purposes.

Parameter	Description	Example Value	Default Value
<code>logonticket_redirect_url</code>	URL that is used to redirect the client to a system that provides SAP logon tickets for SSO authentication	<code>http://link.to.portal/</code> <code>loginService</code>	None

communication

Use the `communication` section of the `xsengine.ini` file to set configuration options for application-related connection requests to SAP HANA, for example, timeouts, port numbers, and maximum number of data end points allowed by the XS engine.

Parameter	Description	Example Value	Default Value
<code>default_read_timeout</code>	Time (in milliseconds) before a connection request is closed	-1, 30, 60	-1 (no time set)
<code>default_read_timeout_override</code>	Ignore setting for <code>default_read_timeout</code>	No, Yes	Yes
<code>listenport</code>	The port number on which the XS Web server listens for requests	30007	3\$(SAPSYSTEM)07

Parameter	Description	Example Value	Default Value
enforced_http_proxy	Override the outgoing proxy settings used for the HTTP/S client, for example, defined in an HTTP/SMTP destination configuration or an <code>httpClient.request()</code> method.	myhost.name.com	None
enforced_https_proxy		myhost.name.com	None
enforced_outbound_proxy	Set the proxy not just for HTTP and HTTPS but for all outgoing protocols, for example: SMTP, socks, ...	myhost.name.com	None
maxchannels	Maximum number of concurrent channels allowed by the XS Web server	Unsigned integer, for example, 1000	4000
maxendpoints	Maximum number of concurrent data endpoints that the XS Web server can expose	Unsigned integer, for example, 1000	4000

customer_usage

The `customer_usage` section of the `xsengine.ini` file is used by the *SAP Web Dispatcher HTTP Tracing* tool to set configuration for SAP HANA application services, for example, to enable HTTP tracing of XS applications on the SAP Web Dispatcher.

Parameter	Description	Example Value	Default Value
<code>/path/to/the/XSapp</code>	The fully qualified path to (and the name of) the application to be traced, for example, <code>sap.hana.ide</code> or <code>sap.hana.xs.admin</code> .	<code>icm/HTTP/logging_n</code>	N/A The parameter <code>/path/to/XSapp</code> is set (or removed) automatically when the administrator uses the <i>XS Admin Tools</i> to enable (or disable) HTTP tracing on the SAP Web Dispatcher for an application.

→ Tip

The parameter value `icm/HTTP/logging_n` is the same as the key defined in the `webdispatcher.ini/profile` section of the configuration parameters, and "n" must be a unique number.

debugger

Use the `debugger` section of the `xsengine.ini` file to set configuration options for the SAP HANA XS JavaScript debugging tools.

Parameter	Description	Example Value	Default Value
enabled	Enable debugging functionality	True/False	False

httpserver

Use the `httpserver` section of the `xsengine.ini` file to set configuration options for the SAP HANA XS Web server, for example, port numbers, and maximum number of sessions and threads allowed.

Parameter	Description	Example Value	Default Value
developer_mode	Enable verbose output for HTTP codes/messages	True/False	False
embedded	Enable the SAP HANA XS engine to run in embedded mode (in the index server). See SAP Note 1849775 .	True/False	False
login_screen_background_image	URL to the image displayed as background in the logon screen, with the following prerequisites: <ul style="list-style-type: none"> File must be reachable by http(s) No requirement for authentication or authorization Recommended minimum resolution of image: 1600*1200 A technical user has to be assigned to the XSSQLCC artifact <code>/sap/hana/xs/selfService/user/selfService.xssqlcc</code>. The technical user must be assigned the role sap.hana.xs.selfService.user.roles.USSExecutor. This user will be used to query the details from the server. 	/sap/hana/xs/ui/Image.jpg	None
max_message_size_mb	Maximum allowed size (in megabytes) of an HTTP request or response	Unsigned integer, for example, 10	100
max_request_runtime	Maximum runtime (in seconds) of an HTTP request targeting an XSJS application. Can be extended in case of long-running database operations.	Unsigned integer, for example, 10	300
maxsessions	Maximum number of registered sessions, not including unauthenticated sessions that are not being debugged	Unsigned integer, for example, 10000	50,000
root_page	Enables requests to the root URI <code>/</code> to be redirected to the URI set with this parameter	/sap/xs/path/root.html	None
sessiontimeout	Amount of time (in seconds) before an inactive session is closed	Unsigned integer, for example, 60	900

odata

Use the `odata` section of the `xsengine.ini` file to set configuration options for OData requests.

Parameter	Description	Example Value	Default Value
<code>allow_nullable_keys</code>	Specify if "key" entity elements can (<code>null</code>) or cannot (<code>not null</code>) have the value <code>NULL</code> .	<code>True/False</code>	<code>False</code>

scheduler

Use the `scheduler` section of the `xsengine.ini` file to set configuration options for the XS job scheduler, which is used to run an XS Javascript or SQLScript as a task in the background at regular intervals

Parameter	Description	Example Value	Default Value
<code>enabled</code>	Activate the XS job-scheduler service. Set to <code>true</code> on one XS host only; this enables <code>xsjob</code> scheduling for the selected instance	<code>True/False</code>	<code>False</code>
<code>sessiontimeout</code>	The amount of time (in seconds) to wait for a job to complete	<code>300 seconds</code>	<code>900 seconds</code>
<code>disable_job_after_restarts</code>	The maximum number of unsuccessful attempts to start a job before the job schedule is automatically disabled	<code>3</code>	<code>5</code>

webdispatcher.ini/profile

Use the `profile` section of the `webdispatcher.ini` file to set configuration options for customer-specific usage scenarios in SAP HANA application services, for example, to enable HTTP tracing of XS applications on the SAP Web Dispatcher.

Parameter	Description	Example Value	Default Value
icm/HTTP/ logging_n	<p>Defines the application-specific log, where “_n” is a unique number. The key's value defines the following:</p> <ul style="list-style-type: none"> • PREFIX= the fully qualified path to (and the name of) the application to be traced, for example, <code>sap.hana.xs.admin</code> • LOGFILE= the location of the log file used to store the trace information; the location includes a variable for the application's name (<code>access_log_app-</code>) and the year, month, and day (<code>%y-%m-%d</code>) • MAXSIZEKB= the maximum allowed size and format of the trace file • SWITCHTF=the time of the day when the new log file is created (DAY/NIGHT) • LOGFORMAT= the format of the trace file content, for example: CLF (common log format), CLFMOD (modified CLF), SAP (SAP log file format), SAPSMD, ... • FLUSH=enable or disable the log flush mechanism 	<pre>PREFIX=/sap/hana/ide/, LOGFILE=\$(_LOCAL_HOST_NAME) / trace/ access_log_sap.hana.ide -%y-%m-%d, MAXSIZEKB=10000 SWITCHTF=day, LOGFORMAT=SAP, FLUSH=1</pre>	<p>N/A</p> <p>The parameter <code>icm/HTTP/logging_n</code> is set (or removed) automatically when the administrator uses the <i>XS Admin Tools</i> to enable (or disable) HTTP tracing on the SAP Web Dispatcher for an application.</p>

→ Tip

The parameter `icm/HTTP/logging_n` is also used as the value for the key defined in the `customer_usage` section of the `xsengine.ini` file.

Related Information

[Default Blacklisted System Properties in Tenant Databases \[page 228\]](#)

12.1.4 Maintaining Application Runtime Configurations

Application runtime configurations specify the security measures that are implemented for access to applications.

The *SAP HANA XS Administration Tool* includes the *XS APPLICATIONS* tool, which you can use to create and maintain runtime configurations for individual applications or a complete application hierarchy. The configuration defined for an application is inherited by any application further down the application package hierarchy. A runtime configuration takes precedence over any runtime configuration located in an application package above it in the package hierarchy.

Note

SAP HANA uses roles to grant access to the features provided by the *SAP HANA XS Administration Tool*. To access the tools required to configure SAP HANA XS runtime configurations, you must have a role based on the role template `sap.hana.xs.admin.roles::RuntimeConfAdministrator` assigned.

You can maintain the following aspects of the application runtime configuration:

→ Tip

Runtime configuration settings override any settings in the application's corresponding application-access (`.xsaccess`) configuration file.

- Application security
Enable/Disable user-authentication checks when starting an application
- User authentication methods (SAML, SPNego, X509, logon tickets, ...)
Enable one or more authentication methods that applications use to authenticate user requests for content.
- Cross Origin Request Sharing (CORS)
Enable support for cross-origin requests, for example, by allowing the modification of the request header. Allowing the sharing of cross-origin resources permits Web pages to make HTTP requests to another domain, where normally such requests would automatically be refused by the Web browser's security policy.
- Custom Headers
Enable support for the X-Frame-Options HTTP header field, which allows the server to instruct the client browser whether or not to display transmitted content in frames that are part of other Web pages. You can also enable this setting in the application's corresponding `.xsaccess` configuration file.
- SQL connection configurations (SQLCC)
Edit the details of an SQL connection configuration, which you use to enable the execution of SQL statements from inside your server-side JavaScript application with credentials that are different to the credentials of the requesting user
- HTTP destination configurations
Edit the details of an HTTP destination configuration, which you use to defines connection details for services running on a specific host, whose details you want to define and distribute
- XS Job Schedules

Edit the details of an XS Job, for example to set the user account under which the job runs, define a start or stop time, and browse the job's log files.

Related Information

[Create an Application Runtime Configuration \[page 1539\]](#)

[Edit and SQL Connection Configuration \[page 1543\]](#)

[Edit an HTTP Destination \[page 1546\]](#)

[Maintain XS Job Details \[page 1619\]](#)

12.1.4.1 Create an Application Runtime Configuration

For SAP HANA XS applications, the runtime configuration defines the security and authentication settings to use when granting access to an application of the content it exposes.

Prerequisites

SAP HANA uses roles to determine the level of access to the features provided by the [SAP HANA XS Administration Tool](#). For example, to access the tools required to perform any tasks relating to application runtime configuration, you must have a role based on the role template `sap.hana.xs.admin.roles::RuntimeConfAdministrator`.

Context

The runtime configuration for an SAP HANA XS application specifies the settings the application uses when it is launched, for example, in response to a user request. If the same settings you define in a runtime configuration are also defined in a design-time file but with a different value, the runtime configuration takes precedence. To create a runtime configuration for an SAP HANA XS application, perform the following steps:

Procedure

1. Start the [SAP HANA XS Administration Tool](#).

The [SAP HANA XS Administration Tool](#) tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who

logs on must have the privileges required to perform administration tasks with the *XS Administration Tools*.

2. Start the SAP HANA *XS Artifact Administration* tool.
In the *Runtime Configuration Details* page you can maintain details of the runtime configurations for the various applications in your package hierarchy.
3. Define the runtime configuration for your application.

i Note

The runtime configuration you define is inherited by all sub-packages in the package hierarchy.

- a. Configure the application security:

Choose the *Edit* button in the *Security & Authentication* tab to can configure the following security options:

- *Public (no authentication required)*
Enable/Disable authentication for application requests.

→ Tip

If you **disable** authentication in the *Security and Authentication* panel, the *Authentication Types* options (for example, SAML or logon tickets) are hidden.

- *Force SSL*
Enable the force SSL option if you want the application to refuse browser requests that do not use secure HTTP (SSL/HTTPS) for client connections.

i Note

The setting for this runtime option overrides the design-time setting for the *force_ssl* keyword in the application's `.xsaccess` file.

- *Prevent Public Access for Sub-Packages*
Ensure that public access only applies to the current package; all subpackages are hidden.

i Note

This option is not available for packages shipped with SAP HANA.

- b. Configure the methods the applications must use to authenticate users.

The *Authentication Types* list is only visible if the *Public (no authentication required)* option is disabled.

i Note

You can select multiple authentication methods which are used in a specific order of priority, for example: first SAML, then logon tickets, and if the user-logon fails for both methods, then basic logon is offered.

To ensure that, during the authentication process, the password is transmitted in encrypted form, it is strongly recommended to enable SSL/HTTPS for all application connections to the XS engine.

- c. Enable support for cross-origin request sharing (CORS), if required.

The *CORS* tab enables you to allow the sharing of cross-origin resources; this permits Web pages to make HTTP requests to another domain, where normally such requests would automatically be refused by the Web browser's security policy.

- d. Enable support for custom headers, if required.

Use the *Custom Headers* tab to configure support for custom headers in the response. This feature enables you to set X-Frame options that allow frames in a Web page to display content from another Web site.

Check the option *Enable Custom Headers* and choose the one of the entries in the list of *X-Frame Options*, for example:

- DENY
- SAMEORIGIN
- ALLOW-FROM <URL>

You can only specify one URL with the ALLOW-FROM option, for example: "value": "ALLOW-FROM http://www.site.com".

i Note

To allow an application to use custom headers, you must enable the *Custom Headers* option.

4. Save the runtime configuration.

i Note

Use the *Reset* button to reset the runtime configuration to its previous state; use the *Revert* button to undo changes to the runtime-configuration options in the current tab.

Related Information

[Application Runtime Configuration Details \[page 1541\]](#)

[Configure HTTPS \(SSL\) for Client Application Access \[page 1575\]](#)

12.1.4.1.1 Application Runtime Configuration Details

In the *XS Artifact Administration* tool, the *Runtime Configuration Details* tab displays information about runtime settings configured for the currently selected application or artifact. You can use the *Runtime Configuration Details* tab to maintain the following details of the runtime configuration:

- [Security & Authentication \[page 1541\]](#)
- [CORS \[page 1542\]](#)
- [Custom Headers \[page 1543\]](#)

Security & Authentication

The *Security & Authentication* tab in the *Runtime Configuration Details* tool enables you to view details of the security settings defined to control access to an application service running on SAP HANA, for example, the

type of access allowed (user/public) and the method used to authenticate users. The following table indicates which information can be defined.

Security and Authentication Details

UI Element	Description	Example
Authentication Type	Enables/Disables requirement for user authentication to access an application service. If you enable authentication, you must select the methods that the application applies to authenticate users, for example, SAML or logon tickets.	Public (No Authentication Required)
Connection Security	Allows only secure HTTPS access to an application; insecure standard HTTP requests are refused. To ensure that passwords are transmitted in encrypted form during the authentication process, it is strongly recommended to enable SSL/HTTPS for all application connections to the XS engine. If you set the <i>force_ssl</i> option, you must ensure that the SAP Web Dispatcher is configured to accept and manage HTTPS requests.	SSL Enforced
Public Access for Sub-Packages	Enables public access to sub packages in an application package hierarchy. This setting cannot be changed for packages shipped with SAP HANA.	Allowed
Authentication Methods	Defines one of more methods that the application service uses to authenticate users requesting access. If multiple methods are selected, an order of priority applies: from most to least secure, for example, <i>SAML</i> , <i>Form Based</i> , and then <i>Basic</i> .	SAML, X509
SAML Identity Provider	The name of the SAML IDP used to verify SAML certificates; this setting is only required if SAML is chosen as one of the authentication methods. <i>Not Applicable</i> indicates that no SAML IDP is configured.	SAMLIDP1

CORS

The *CORS* tab in the *Runtime Configuration Details* tool enables you to view details of the settings defined to control access to your application resource from other Web browsers. For example, you can specify where requests can originate from or what is allowed in the request and response headers. The following table indicates which information can be defined for Cross Origin Resource Sharing.

CORS Settings

CORS Option	Description
Cross Origin Resource Sharing	Enable/Disable requests from other browser sessions to an application.
ALLOWED ORIGINS	A single host name or a comma-separated list of host names that are allowed by the server, for example: <code>www.sap.com</code> or <code>*.sap.com</code> . If no host is specified, the default <code>**</code> (all) applies. Note that matching is case-sensitive.
ALLOWED HEADERS	A single header or a comma-separated list of request headers that are allowed by the server. If no request header is specified, no default value is supplied.

CORS Option	Description
EXPOSED HEADERS	A single header or a comma-separated list of response headers that are allowed to be exposed. If no response header is specified for exposure, no default value is supplied.
ALLOWED METHODS	A single permitted method or a comma-separated list of methods that are allowed by the server, for example, "GET", "POST". If no method is specified, the default "GET", "POST", "HEAD", "OPTIONS" (all) applies. Note that matching is case-sensitive.
MAX AGE	A single value specifying how long a preflight request should be cached for. If no value is specified, the default time of "3600" (seconds) applies.

Custom Headers

The [Custom Headers](#) tab in the [Runtime Configuration Details](#) tool enables you to configure support for custom headers in the HTTP response. This feature enables you to set X-Frame options that allow frames in a Web page to display content from another Web site.

Custom Headers Details

UI Element	Description	Example
Custom Headers	Enable/Disable the use of custom headers in HTTP response.	Disabled
X-Frame Options	Allow/Deny requests to display content from the same or another Web site. Note that you can only specify one URL with the ALLOW-FROM option, for example: "value": "ALLOW-FROM http://www.site.com".	DENY, SAMEORIGIN, ALLOW-FROM

Related Information

[Create an Application Runtime Configuration \[page 1539\]](#)

12.1.4.2 Edit an SQL Connection Configuration

In SAP HANA Extended Application Services (SAP HANA XS), you use the SQL connection configuration to enable the execution of SQL statements from inside your server-side JavaScript application with credentials that are different to the credentials of the requesting user.

Prerequisites

SAP HANA uses roles to determine the level of access to the features provided by the [SAP HANA XS Administration Tool](#). For example, to access the tools required to perform any tasks relating to SQL connection

configuration (SQLCC), you must have a role based on the role template `sap.hana.xs.admin.roles::SQLCCAdministrator`. This role includes the related role `sap.hana.xs.admin.roles::SQLCCViewer`

Context

The SQL connection configuration enables the execution of SQL statements from inside your server-side JavaScript application with credentials that are different to the credentials of the requesting user. You can use the *XS Artifact Administration* tool to change the user name in the XS SQL connection-configuration file.

Procedure

1. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must have the privileges required to perform administration tasks with the *XS Artifact Administration* tool.

2. Start the SAP HANA *XS Artifact Administration* tool.

In the *XS Artifact Administration* tool you can manage the runtime configurations for the various applications in your package hierarchy.

3. Locate the SQL connection configuration object.

In the *Application Objects* list, locate and select the object containing the SQL connection configuration that you want to edit; SQL connection configuration objects have the file extension `.xssqlcc`. The details are displayed in the *SQL Connection Details* panel.

4. Maintain the SQL connection details.

The *SQL Connection Details* allows you to modify details of the database user whose credentials are used to establish the SQL connection defined in the SQLCC object. If a role is specified in the `role_for_auto_user` parameter, SAP HANA assigns the role defined in `role_for_auto_user` to the new auto-generated user.

i Note

If you bind the XS SQL connection to a specific existing database user (not the auto user), you must provide the user's password. If do not provide a password for the specified database user, you cannot save the changes to the SQLCC object's runtime configuration.

5. Set the run-time status of the XS SQL connection configuration.

You must set the runtime status of the XS SQL connection configuration to *Active*; the run-time status can only be changed by an administrator. When the run-time status of the XSQL connection configuration is

set to *active*, SAP HANA automatically generates a new user (XSSQLCC_AUTO_USER_[...]) for the XSSQL connection configuration object and assigns the role defined in `role_for_auto_user` to the new auto-generated user.

6. Save the changes.

Related Information

[SQL Connection Details \[page 1545\]](#)

12.1.4.2.1 SQL Connection Details

The SQL-connection configuration file specifies the details of a connection to the database.

The database connection established by the SQL-connection configuration file enables the execution of SQL statements with credentials that are different to the credentials of the requesting user, for example, from inside a server-side (XS) JavaScript application.

The *SQL Connection Details* tab in the *XS Artifact Administration* tool enables you to view details of the XS SQL connection configurations that you have defined, for example, the package location, and the user bound to the SQL connection. The following table indicates which information can be viewed.

SQL Connection Details

UI Element	Description	Example
<i>Package</i>	The name of the repository package containing the currently selected SQL connection configuration	testApp
<i>Description</i>	A short description of the selected SQL connection configuration	Admin SQL connection
<i>Username</i>	The name of the user to whom you want to bind the SQL connection configuration. If no user is specified, SAP HANA automatically generates the user XSSQLCC_AUTO_USER_[...] when the run-time status of the XSSQL connection configuration is set to <i>Active</i> . The new auto-user is assigned the role specified in <i>Role for Auto User</i> . If you bind the SQL connection manually to a specific SAP HANA user, you must supply the user's password to enable a connection to be established and ensure that the user has the necessary privileges (for example, by assigning a role).	XSSQLCC_AUTO_USER[...]
<i>Password</i>	The password for the user bound to the SQL connection configuration. A password is not required for the automatically generated XSSQLCC_AUTO_USER_[...].	*****
<i>Assigned by</i>	The name of the user who assigned the user defined in <i>Username</i> to the currently selected SQL connection configuration	JohnDoe

UI Element	Description	Example
<i>Role for Auto User</i>	The name of (and package path to) the role to be assigned to the new auto-user that is generated when the run-time status of the XSSQL connection configuration is set to <i>active</i>	acme.com.xs.roles::JobAdministrator
<i>Status</i>	The current runtime status of the XSSQL connection configuration (active/inactive)	active

Related Information

[Edit an SQL Connection Configuration \[page 1543\]](#)

12.1.4.3 Edit an HTTP Destination Runtime Configuration

An HTTP destination defines connection details for services running on a specific host, whose details you want to define and distribute. The HTTP destination can be referenced by an application.

Prerequisites

SAP HANA uses roles to determine the level of access to the features provided by the [SAP HANA XS Administration Tool](#). For example, to access the tools required to perform any tasks relating to HTTP destination configuration (HTTPDest), you must have a role based on the role template `sap.hana.xs.admin.roles::HTTPDestAdministrator`.

Context

To edit an HTTP destination using the [SAP HANA XS Administration Tool](#), perform the following steps:

Procedure

1. Start the [SAP HANA XS Administration Tool](#).

The [SAP HANA XS Administration Tool](#) tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who

logs on must have the privileges required to perform administration tasks with the *XS Artifact Administration* tool.

2. Start the SAP HANA *XS Artifact Administration* tool.
Use the *XS Artifact Administration* tool to manage the runtime configurations for the various applications in your package hierarchy.
3. Locate the HTTP destination configuration object that you want to edit.
In the *Application Objects* list, locate and select the object containing the HTTP destination configuration that you want to edit. HTTP destination configuration objects have the suffix `.xshttpdest`. The details are displayed in the *HTTP Destination Details* panel.
4. Edit the details of the HTTP destination configuration.

To edit the details of an HTTP destination configuration, choose the *Edit* button in the screen displaying the details you want to edit, for example:

- *General Information*
Host name and port of the server to connect to, any path prefix (to add to the start of the URL used to connect to the service on the remote host, and a timeout setting for the time allowed to connect to the remote host.

i Note

The *Extends* option is only available if the HTTP destination you are modifying is being used to extend the configuration defined in another HTTP destination.

- *Proxy Details*
Details of the proxy type (*None*, *HTTP*, or *Socks*), the name of the system hosting the proxy service, the port to connect on and the user credentials required to establish the connection.

⚠ Caution

The proxy-server settings you define here are overridden by any SAP HANA system wide setting for a proxy server, for example, defined by the `enforced_outbound_proxy` parameter in the `communication` section of the `xsengine.ini` configuration file.

- *Authentication Details*
 - *SSL Enabled*
SSL for outbound connections between SAP HANA XS and the host named in the HTTP destination configuration.
You must choose an *SSL Authentication Type*. If you choose *Client Certificate* (default), you must specify the *Trust Store* where the certificates are stored. You can choose an existing trust store from a list of stores configured for the SAP HANA instance (in the *Trust Store* drop-down menu), or create a new trust store using the *Trust Manager*.
SSL Host Check (`true | false`) enables a check which verifies that the certificate used for authentication is valid (matches the host). If the certificate does not match, SSL terminates.
 - *Authentication Type* (for example, *none*, *basic*, *SAP Assertion Ticket*, *SAML*, or *SAML Assertion Propagation*).

i Note

The *Authentication Type* you select determines what (and how much) additional information is required.

For example, for the [SAP Assertion Ticket](#) authentication type, you must provide the SAP SID and client number of the instance providing the service. The value displayed (if any) is the one already defined in the design-time representation of the HTTP destination configuration. Any changes you make to the runtime configuration (here) are synchronized with the design-time configuration artifact.

For [SAML](#), the values displayed reflect the parameters set in the corresponding design-time representation of the HTTP destination configuration, for example,

`ConfigFileName.xshttpdest`. For more information, see [HTTP Destination Details](#) in [Related Information](#).

- [OAuth Details](#)

You cannot enter this information manually; the information is read from the design-time configuration file that describes the OAuth application, for example, `oauthDriveApp.xsoauthappconfig`. To display a list of available OAuth application-configuration packages (files with the suffix `*.xsoauthappconfig`) on your SAP HANA system, choose [Browse OAuth App Configs](#) and select a package from the list. The location of the package containing the OAuth application-configuration you choose is used to populate the [OAuth App Config Package](#) field; the name of the OAuth application-configuration you choose is used to populate the [OAuth App Config Name](#) field.

5. Save the changes.

Saving the changes to the HTTP destination configuration automatically commits the HTTP destination configuration object to the SAP HANA repository and activates it.

→ Tip

Use the [Reset](#) button to reset the runtime configuration to its previous state; use the [Revert](#) button to undo changes to the runtime-configuration options in the current tab.

Related Information

[HTTP Destination Details \[page 1548\]](#)

[SAP HANA XS Classic Classic Configuration Parameters \[page 1532\]](#)

12.1.4.3.1 HTTP Destination Details

An HTTP destination defines connection details for services running on a specific host, whose details you want to define and distribute

In the [XS Artifact Administration](#) tool, the [HTTP Destination Details](#) tab displays information about the currently selected HTTP destination. You can use the [HTTP Destination Details](#) tab to maintain the following details of the runtime configuration:

- [General Information \[page 1549\]](#)
- [Proxy Details \[page 1549\]](#)
- [Authentication Details \[page 1550\]](#)
- [OAuth Details \[page 1551\]](#)

General Information

The *General Information* tab in the *HTTP Destination Details* tool enables you to view details of the HTTP destination that you have defined, for example, the name of the destination host, the port to connect on, and a short description. The following table indicates which information can be viewed.

HTTP Destination Details

UI Element	Description	Example
<i>Extends</i>	The name of another HTTP destination configuration which the currently selected configuration is using as a base but also modifying.	gfn.xshttpdest
<i>Description</i>	A short description of the selected HTTP destination	Service @ Destination
<i>Host</i>	The name of the system hosting the services defined in the HTTP destination configuration	download.finance.acme.com
<i>Port</i>	The port to connect to on the remote host	80
<i>Path Prefix</i>	The prefix to add to the start of the URL used to connect to the service on the remote host	/d/quotes.csv?f=a
<i>Timeout</i>	The time allowed to connect to the remote host defined in the HTTP destination	0

Proxy Details

The *Proxy Details* tab in the *HTTP Destination Details* tool enables you to view details of the proxy service used by the HTTP destination that you have defined, for example, the name of the proxy host, the port to connect on, and the user credentials required to establish a connection. The following table indicates which information can be viewed and configured.

Proxy Server Details

UI Element	Description	Example
<i>Proxy Type</i>	The type of proxy service, for example: None, HTTP, or SOCKS.	HTTP
<i>Proxy Host</i>	The name of the system hosting the proxy service used by the HTTP destination	proxy.host.acme.com
<i>Proxy Port</i>	The port to connect to on the system hosting the proxy service	8080
<i>Proxy User</i>	The user credentials required to connect to the proxy service	johndoe

Note

The proxy-server settings you define here are overridden by any SAP HANA system-wide setting for a proxy server, for example, defined by the `enforced_outbound_proxy` parameter in the `communication` section of the `xsengine.ini` configuration file.

Authentication Details

The *Authentication Details* tab in the *HTTP Destination Details* tool enables you to view details of the authentication service used by the HTTP destination that you have defined, for example, the authentication **type** and the trust store used to maintain any SSL client certificates. The following table indicates which information can be viewed and modified.

Authentication Details

UI Element	Description	Example
<i>Authentication Type</i>	The type of service used for authentication, for example: <i>None</i> , <i>Basic</i> , <i>SAP Assertion Ticket</i> , <i>SAML</i> , or <i>SAML Assertion Propagation</i>	<i>SAML</i>
<i>Communication Security</i>	Enable or disable SSL communication. If you enable SSL, you must select an <i>SSL Authentication Type</i> .	<i>SSL Enabled</i>
<i>SSL Authentication Type</i>	The type of authentication used for SSL, for example, <i>Client Certificate</i> (default) or <i>Anonymous</i> . If you choose <i>Client Certificate</i> , you must specify the trust store where the certificates are located.	<i>Client Certificate</i>
<i>SSL Host Check</i>	Enable or disable the SSL host check; the check verifies that the certificate used for authentication is valid (matches the host).	Enabled
<i>Trust Store</i>	The name of the trust store used to maintain security certificates required during the authentication process; select from the drop-down list	SAPLogon

The following table lists the choices available when configuring the authentication type for an HTTP destination.

HTTP Destination: Authentication Type

UI Element	Description	Example
<i>None</i>	No user authentication is performed	-
<i>Basic</i>	The <i>Name</i> of the user whose account is used to log on to the HTTP destination using basic authentication	JohnDoe
	The <i>password</i> of the user specified in <i>Name</i>	*****
<i>SAP Assertion Ticket</i>	System ID (<i>SAP SID</i>) of the SAP instance providing the SAP Assertion Ticket service	GFN
	<i>Client number</i> of the SAP instance providing the SAP Assertion Ticket service	007
<i>SAML</i>	The <i>Entity ID</i> of the remote SAML party	accounts.acme.com
	<i>User Mapping</i> : a list of name-ID mappings, for example, <i>Unspecified</i> , <i>Email</i> , <i>Email</i> , <i>Unspecified</i>	Email
	<i>Assertion Consumer Service</i> defines the way in which SAML assertions and responses are sent, for example: as an authorization header or POST parameter.	Assertion as POST parameter
	Additional <i>Attributes</i> for the SAML Assertion.	Email

UI Element	Description	Example
SAML Assertion Propagation	Allow an SAML token to be forwarded from the server where the token was received to another server.	N/A

For the authentication type [SAML](#), the values displayed reflect the parameters set in the corresponding design-time representation of the HTTP destination configuration, as illustrated in the following table:

HTTP Destination SAML Runtime:Design-Time Parameters

SAML Runtime Parameter	SAML Design-Time Parameter	Description
Entity ID	samlProvider	The entity ID of the remote SAML party
Assertion Consumer Service	samlACS	The way in which SAML assertions or responses are sent
Attributes	samlAttributes	Additional attributes for the SAML Assertion.
User Mapping	samlNameId	A list of name-ID mappings, for example, e-mail .

OAuth Details

The [OAuth Details](#) tab in the [HTTP Destination Details](#) tool enables you to view details of the OAuth package used by the HTTP destination that you have defined. An OAuth configuration package is a collection of configuration files that define the details of how an application uses OAuth to enable logon to a resource running on a remote HTTP destination. The following table indicates the information that can be viewed.

HTTP Destination: OAuth Information

UI Element	Description	Example
OAuth App Config Package	The name of the repository package containing the OAuth application-configuration	sap.hana.xs.oauth.lib.providerconfig
OAuth App Config Name	The name of the OAuth application-configuration (repository artifacts with the suffix <code>.xsoauthappconfig</code>)	abap.xsoauthappconfig

Note

You cannot enter this information manually; the information is read from the design-time configuration file that describes the OAuth application, for example, `oauthDriveApp.xsoauthappconfig`.

Related Information

[Edit an HTTP Destination Runtime Configuration \[page 1546\]](#)

12.1.5 Managing Trust Relationships

Trust relationships enable you to establish secure connections between known servers whose identity can be confirmed by a signed certificate. The certificates are stored in a trust store.

The *SAP HANA XS Administration Tool* includes the *Trust Manager*, which is an application that you can use to create and maintain the certificates used to establish trust relationships between servers. You can use the *Trust Manager* to perform the following tasks.

i Note

SAP HANA uses roles to grant access to the features provided by the *SAP HANA XS Administration Tool*. To access the tools required to maintain trust relationships between SAP HANA and other systems, you must have a role based on the role template `sap.hana.xs.admin.roles::TrustStoreAdministrator`.

- Add/Delete a trust store
SAP HANA makes use of multiple trust stores. The trust stores listed below are provided by default.

i Note

The trust stores listed below are located in the **file system**. In some cases, it is possible and recommended to use trust stores that exist in the database as database objects. In-database trust stores (referred to as certificate collections) contain the required client certificates, which are also stored in the database. We recommend using in-database certificate collections where possible. For more information, see *Managing Client Certificates in the SAP HANA Database*.

- The SAP HANA trust store (`sapshr.pse`)
Used for secure SQL and SAML or OAuth scenario, `sapshr.pse` is installed automatically and is available by default.

→ Recommendation

For user authentication based on X.509 certificates and SAML assertions, we recommend creating separate certificate collections with the purposes *X.509* and *SAML* instead of using the file system-based trust store `sapshr.pse`.

- The SAP Web Dispatcher trust store (`SAPSSLS.pse`)
Required for SSL connections using the Secure Socket Layer, `SAPSSLS.pse` is installed automatically and is available by default.
- The SAP Logon Ticket trust store (`saplogon.pse`)
Optional: `saplogon.pse` is only necessary if an SAP HANA XS application requires an SAP logon ticket from a user at logon

→ Recommendation

For user authentication based on logon tickets, we recommend creating a certificate collection with the purpose *SAP LOGON* instead of using the file system-based trust store `saplogon.pse`.

- The client authentication trust store (`SAPSSLC.pse`)
Optional: `SAPSSLC.pse` is only required for client connections, for example, that use the SQL client interface (`hdbsql`).
- Manage your own certificates

- Import a private key
- Create a certificate request
- Have the requested certificate signed by a certificate authority
- Import the signed certificate into the trust store
- Manage server certificates
 - HTTP destinations (via SSL/HTTPS)
 - Certificate authorities (for example, “Verisign” or “TC TrustCenter Universal”)

The *Trust Manager* tool enables you to configure the out-bound view; that is, trust relationships with remote systems that provide services required by SAP HANA XS applications. If you want to configure the **in-bound** view (for example, incoming requests **to** SAP HANA), use the SAP HANA *Web Dispatcher Administration* tool.

- Out-bound trust
Secure communication and trust for out-bound communication, for example, between an SAP HANA XS application and an ABAP system using using SSL/TLS.
- In-bound trust
Secure communication and trust for in-bound communication, for example, between an SAP HANA XS application and an ABAP system using using SSL/TLS.

Both the *Trust Manager* and the *Web Dispatcher Administration* tools are available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/wdisp/admin`.

i Note

Access to the *Web Dispatcher Administration* tool is enabled by the role `sap.hana.xs.wdisp.admin::WebDispatcherAdmin`.

Related Information

[Add/Edit a Trust Store \[page 1554\]](#)

[Import a Server Certificate \[page 1557\]](#)

[Create Your Own Certificate \[page 1555\]](#)

12.1.5.1 Add/Edit a Trust Store

The trust store enables you to maintain a list of servers that you trust; the trust is based on a certificate you import into the trust store and which can be signed by a certificate authority, for example, Verisign or TCTrustCenter.

Prerequisites

SAP HANA uses roles to determine the level of access to the features provided by the SAP HANA XS Administration Tool. To access the tools required to add a trust store, you must have a role based on the role template `sap.hana.xs.admin.roles::TrustStoreAdministrator`.

Context

→ Recommendation

This procedure describes how to create a trust store in the file system. We recommend creating trust stores in the database (referred to as certificate stores) where possible. For more information, see the section *Managing Client Certificates in the SAP HANA Database*.

To enter the details of trust store, you can use the *SAP HANA XS Administration Tool*, as described in the following steps.

⚠ Caution

To maintain the details of a trust store, you must be familiar with the concepts of trust stores and the certificates they contain.

Procedure

1. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must also have the privileges required to perform the administration tasks associated with trust stores.

2. Start the *Trust Manager*.

The *Trust Manager* is available in the list of SAP HANA XS administration tools.

3. Create the new trust store.

In the *Create Trust Store* dialog, you must provide a name for the new trust store.

- a. In the *Trust Store* pane, choose *Add* to open the *Create Trust Store* dialog.
- b. Type a name for the new trust store and choose *OK*.
Choose *OK* to add the trust store to the list of trust stores known to SAP HANA XS.

4. Define the details of the new trust store.

You can use the *Own Certificate* and *Certificate List* to manage the certificates you import for the servers that are known to and trusted by SAP HANA XS.

Related Information

[Import a Server Certificate \[page 1557\]](#)

[Create Your Own Certificate \[page 1555\]](#)

[Managing Client Certificates \[page 900\]](#)

12.1.5.2 Create Your Own Certificate

The trust store enables you to maintain a list of servers that you trust; the trust is based on a certificate you import into the trust store and which can be signed by a certificate authority, for example, Verisign or TCTrustCenter.

Prerequisites

i Note

This feature is available with restricted releases. If you want to use it, refer to SAP Note 1779803. See the Related Information section for the direct link.

SAP HANA uses roles to determine the level of access to the features provided by the *SAP HANA XS Administration Tool*. To access the tools required to perform trust manager tasks, you must have a role based on the role template `sap.hana.xs.admin.roles::TrustStoreAdministrator`.

Context

You can use the certificates stored in the trust store to secure the communication between trusted servers, for example, with SSL/HTTPS. However, you must also create a certificate that you can use to authenticate the identity of the SAP HANA server, too.

To create your own certificate and import it into your trust store, perform the following steps:

Procedure

1. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

Note

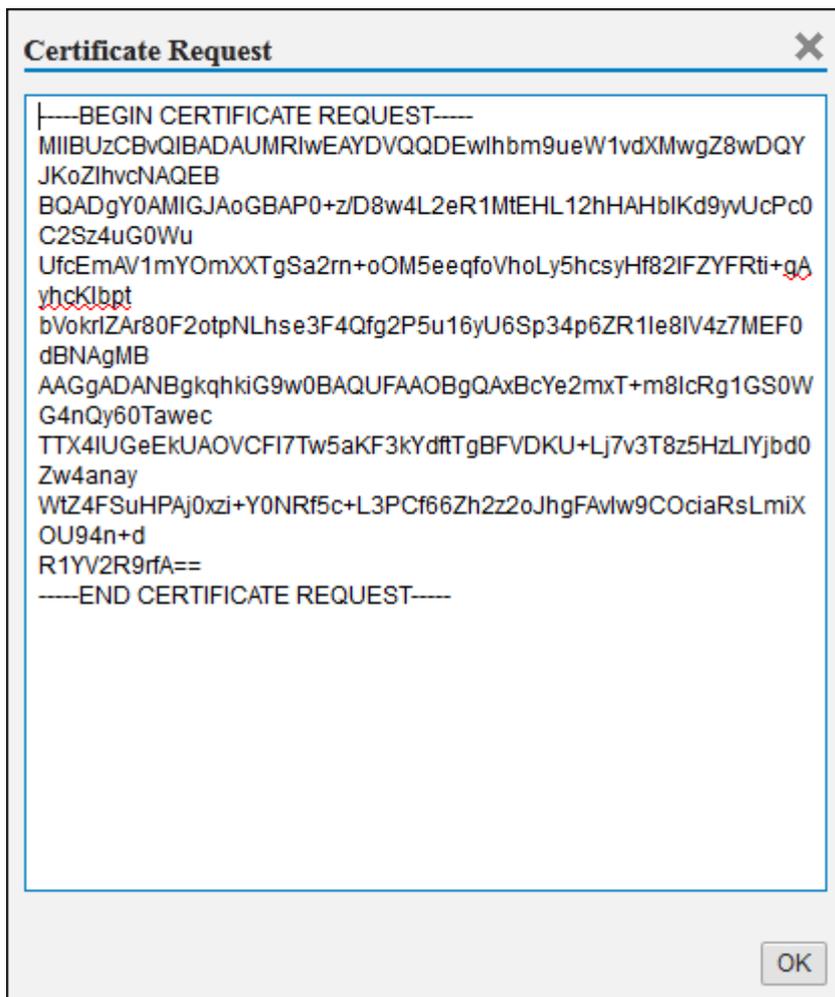
In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must have the privileges required to perform administration tasks with the *Trust Manager* tool.

2. Start the SAP HANA XS *Trust Manager* tool.

In the list of tools, choose *Trust Manager* tab to display the screen where you can manage the certificates in your trust store.

3. Create a certificate request.

In the *Own Certificate* panel, choose **► Certificate Actions ► Create CA Request ►**.



4. Send the certificate request to a certificate authority for signing.

You must send the certificate request to a certificate authority (CA) to have it signed; you import the response from the CA into your trust store.

5. Import the signed certificate into the trust store.

This may be a trust store in the file system (for example, `sapsrv.pse`) or an in-memory certificate collection with the purpose *SAML* (recommended).

Option	Description
Certificate collection with purpose <i>SAML</i> (recommended)	Use the SAP HANA cockpit to import the certificate into the certificate store and then add it to the relevant collection. For more information, see the section on managing certificates.
Trust store in the file system	In the <i>Own Certificate</i> panel, choose ► <i>Certificate Actions</i> ► <i>Put CA Response</i> ►. The imported certificate is displayed in the certificate list.

Related Information

[SAP Note 1779803](#)

[Add/Edit a Trust Store \[page 1554\]](#)

[Managing Client Certificates \[page 900\]](#)

12.1.5.3 Import a Server Certificate

A server certificates enables you to establish a trusted relationship between SAP HANA and the server described in the server certificate. You import the certificates into the trust store.

Prerequisites

SAP HANA uses roles to determine the level of access to the features provided by the *SAP HANA XS Administration Tool*. To access the tools required to perform trust manager tasks, you must have a role based on the role template `sap.hana.xs.admin.roles::TrustStoreAdministrator`.

Context

→ Recommendation

This procedure describes how to import a server certificate into a trust store in the file system. We recommend creating trust stores in the database (referred to as certificate stores) where possible. For

more information about how to import certificates into the in-memory certificate store and add them to certificate collections, see the section *Managing Client Certificates in the SAP HANA Database*.

The trust store enables you to maintain a list of servers that you trust; the trust is based on a certificate you import into the trust store and which can be signed by a certificate authority, for example, Verisign or TCTrustCenter. You can use the certificates to secure the communication between the trusted servers, for example, with SSL/HTTPS.

To import a certificate into your trust store, perform the following steps.

Procedure

1. Obtain a copy of the certificate you want to import into your trust store.
You can export a certificate from a server and save it to a temporary location.

2. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must have the privileges required to perform administration tasks with the *Trust Manager* tool.

3. Start the SAP HANA XS *Trust Manager* tool.
In the list of tools, choose the *Trust Manager* tab to display the screen where you can manage the certificates in your trust store.

4. Locate the copy of the certificate you want to import into the trust store.

In the *Certificate List* panel, choose **▶ Import Certificate ▶ Browse... ▶** and navigate to the folder containing the certificate you want to import.

i Note

Trust certificates usually have a recognizable suffix such as `.crt`, for example, `TCTrustCenterUniversalCAIII.crt`.

5. Import the certificate into the trust store.

In the *Import Certificate* dialog, choose *Import Certificate*.

i Note

If you are importing a certificate you created yourself, you must provide a password to complete the import operation.

The imported certificate is displayed in the certificate list.

Related Information

[Add/Edit a Trust Store \[page 1554\]](#)

[Managing Client Certificates \[page 900\]](#)

12.1.6 Maintaining SAML Providers

You can configure an SAP HANA system to act as a service provider for Single Sign On (SSO) authentication based on Security Assertion Markup Language (SAML) certificates.

The *SAP HANA XS Administration Tool* includes the *SAML CONFIGURATION* application, which you can use to configure SAP HANA system to act as an SAML service provider for SSO authentication. You must perform this step if you want your SAP HANA XS applications to use SAML as the logon authentication method, for example, by enabling the *SAML* option in the *AUTHENTICATION* panel in the *XS APPLICATIONS* tool

i Note

SAP HANA uses roles to grant access to the features provided by the *SAP HANA XS Administration Tool*. To access the tools required to configure an SAP HANA system to act as an SAML service provider, you must have a role based on the role template `sap.hana.xs.admin.roles::SAMLAdministrator`.

You can use the *SAML CONFIGURATION* tool to perform the following tasks:

- Configure an SAP HANA system to act as a service provider
- Add a new SAML Identity provider (IDP)
- Modify the details of an existing SAML Identity provider (IDP)

i Note

To maintain a SAML identity provider (IDP), you must be logged on to SAP HANA with the credentials of the system user.

Related Information

[Configure an SAP HANA System as an SAML Service Provider \[page 1560\]](#)

[Add a SAML Identity Provider in SAP HANA Studio \[page 732\]](#)

12.1.6.1 Configure an SAP HANA System as an SAML Service Provider

SAP HANA supports the use of authentication based on Security Assertion Markup Language (SAML) certificates.

Prerequisites

SAP HANA user roles are used to determine the level of access to the features provided by the [SAP HANA XS Administration Tool](#). To access the tools required to configure an SAP HANA system to act as an SAML service provider, you must have a role based on the role template `sap.hana.xs.admin.roles::SAMLAdministrator`.

Context

You can configure an SAP HANA system to act as a service provider for authentication based on Security Assertion Markup Language (SAML) certificates. You must perform this step if you want your the SAP HANA XS applications to use SAML as the user authentication method.

⚠ Caution

To maintain the details of SAML service providers, you must be familiar with the technical background of SAML SSO mechanisms and requirements.

Procedure

1. Start the [SAP HANA XS Administration Tool](#).

The [SAP HANA XS Administration Tool](#) tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must have the privileges required to perform SAML administration tasks.

2. Start the SAP HANA XS [SAML Service Provider](#) tool.
In the list of tools, choose [SAML Service Provider](#) to display the screen where you can enter details of the SAML service provider you want to configure.
3. Enter details of the SAML service provider.
In the [Service Provider Information](#) panel choose [Edit](#); you must provide the following information:

i Note

The information you enter is used to populate the XML document saved as the SAML service-provider metadata.

- **Name**
This can be any name but, for troubleshooting purposes, is usually the fully qualified name of the system hosting the SAML service.
- **Organization Name**
According to the oasis SAML standard, the name of the organization responsible for the SAML service described here. The name you enter here is wrapped in the XML tag <OrganizationName> used in the SAML certificate. The organization name can (but does not have to) be human readable.
- **Organization Display Name**
According to the oasis SAML standard, the human-readable form of the name of the organization responsible for the SAML service described here. The name you enter here is wrapped in the XML tag <OrganizationDisplayname> that is contained in the SAML certificate.
- **Organization URL**
A URL that specifies a location where a user can find additional information about the organization responsible for the SAML service you describe in this task.

The information you enter in the various configuration tabs and screens is added to the appropriate tags in the XML document displayed in the *Metadata* tab.

4. Save the SAML service-provider configuration.

Choose *Save*; the XML document describing the SAML service is parsed and, if no errors are found, saved.

Related Information

[SAML Service Provider Details \[page 1561\]](#)

12.1.6.1.1 SAML Service Provider Details

An SAP HANA system can act as an SAML service provider for SSO authentication.

An SAP HANA system can act as a service provider for authentication based on Security Assertion Markup Language (SAML) certificates. The *SAML Service Provider* tool displays the following screens to help you maintain details of the SAML service provider:

- [Service Provider Information \[page 1562\]](#)
- [Service Provider Configuration \[page 1562\]](#)
- [Metadata \[page 1562\]](#)

i Note

The information you enter is used to populate the XML document saved as the SAML service-provider metadata.

Service Provider Information

The *Service Provider Information* tab in the *SAML Service Provider* tool enables you to provide details of the SAML service provider. The following table indicates which information is required.

UI Element	Description	Example
Name	The fully qualified name of the system hosting the SAML service	SAMLSP01
Organisation Name	The name of the organisation responsible for the SAML service provider. The name you enter is wrapped in the XML tag <OrganizationName> used in the SAML certificate. <i>Organization Name</i> can (but does not have to) be human readable	SAP
Organisation Display Name	The human-readable name of the organisation responsible for the SAML service provider. The name you enter here is wrapped in the XML tag <OrganizationDisplayname> used in the SAML certificate.	SAP
Organisation URL	A location where a user can find additional information about the organization responsible for the SAML service	sap.com

Service Provider Configuration

The *Service Provider Configuration* tab in the *SAML Service Provider* tool enables you to maintain details of the SAML service provider used to handle SAML assertions. The following table indicates which information is required.

UI Element	Description	Example
Hash	The hash algorithm use to encode SAML assertions	SHA256
Add Key Info	If <Keyinfo> node should be included in the XML signature; default = yes	"yes" or "no"
Default Application Path	Path to the application requiring logon user credentials provided by the SAML service provider, if the SSO request is initiated by an SAML identity provider	/
Assertion Timeout	Period of time (in seconds) for which SAML assertion requests for SSO initiated by an SAML service provider remain valid; default=10 minutes	1000
Default Role	Default SAP HANA role assigned to new SAML users	JobViewer

Service Provider Metadata

The *Metadata* tab in the *SAML Service Provider* tool enables you to view details of the SAML service provider used to handle SAML assertions. The metadata document includes the information you enter in the *Service Provider Information* and *Service Provider Configuration* tabs.

Field Name	Description
Metadata	An XML file containing details of the SAML service provider used to handle SAML assertions

12.1.6.2 Add an SAML Identity Provider

SAP HANA supports the use of SSO authentication based on Security Assertion Markup Language (SAML) certificates. An identity provider is used by the service provider to authenticate the users signing in by means of SSO.

Prerequisites

- The SAP HANA trust store contains the server certificate that will be used to generate SAML SP metadata and validate SAML assertions (service provider certificate). We recommend that you use an in-memory certificate collection with purpose [SAML](#). For more information, see the section on managing client certificates.
- SAP HANA user roles are used to determine the level of access to the features provided by the SAP HANA XS Administration Tool. To access the tools required to add an SAML identity provider (SAML IDP), you must have a role based on the role template `sap.hana.xs.admin.roles::SAMLAdministrator`.
- You need access to the XML document containing the IDP metadata that describes the SAML identity provider (SAML IDP) you want to add.

Context

To enter the details of SAML identity providers, you can use the [SAP HANA XS Administration Tool](#), as described in the following steps:

⚠ Caution

To maintain the details of an SAML identity provider, you must be familiar with the technical background of SAML SSO mechanisms and requirements.

Procedure

1. Start the [SAP HANA XS Administration Tool](#).

The [SAP HANA XS Administration Tool](#) tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must have the privileges required to perform SAML administration tasks.

2. Add an SAML SSO identity provider (IDP).

The information required to maintain details of an SAML IDP is specified in an XML document containing the IDP metadata. This document should be available as part of the SAML service you want SAP HANA XS to use. The only information you must provide manually is the name of the new IDP; the IDP name must be unique.

- a. In the *SAML Identity Provider List*, choose **[+]** to display the *Add Identity Provider Info* pane.
- b. In the *Add Identity Provider Info* pane, paste the contents of the XML document containing the IDP metadata into the *Metadata* box.

If the contents of the XML document are valid, the parsing process extracts the information required to insert into the *Subject*, *Entity ID*, and *Issuer* fields in the *General Data* screen area, and the URL fields in the *Destination* screen area, for example, *Base URL* and *SingleSignOn URL (*)*.

- c. In the *Name* box of the *General Data* screen area, enter a name for the new SAML SSO identity provider.

i Note

The name of the SAML IDP is mandatory and must be unique; it appears in the list of available SAML IDPs that is displayed, if you select SAML as the authentication method for SAP HANA XS applications to use, for example, in the *Authentication* screen area of the *XS Artifact Administration* tool.

3. Save the details of the new SAML identity provider.

Choose **Save** to save the details of the SAML identity provider and add the new SAML IDP to the list of known SAML IDPs.

The new SAML IDP is displayed in the list of known IDPs shown in the *SAML Identity Provider List*.

4. Check the details of the new SAML IDP.

Select the new SAML IDP in the list of known SAML IDPs to display the IDP's details in the information panel.

Next Steps

Copy the certificate from the SAML IDP metadata document and add it to the SAP HANA trust store for SAML authentication (certificate collection with purpose *SAML*). For more information, see *Configure SSO with SAML Authentication for SAP HANA XS Applications*.

Related Information

[Configure an SAP HANA System as an SAML Service Provider \[page 1560\]](#)

[Modify an Existing SAML Identity Provider \[page 1566\]](#)

[SAML Identity Provider Details \[page 1565\]](#)

[Managing Client Certificates \[page 900\]](#)

[Configure SSO with SAML Authentication for SAP HANA XS Applications \[page 1586\]](#)

12.1.6.2.1 SAML Identity Provider Details

An SAML identity provider is used by the SAML service provider to authenticate users signing in by means of a single sign-on (SSO) mechanism.

SAP HANA supports the use of SSO authentication based on Security Assertion Markup Language (SAML) certificates. An SAML identity provider is used by the SAML service provider to authenticate users who sign in to an application by means of SSO. As part of the SAML IDP configuration, you specify the following options:

- [General data \[page 1565\]](#)
- [HTTP Destination \[page 1565\]](#)

General Data

The *General Data* screen area in the *SAML Identity Provider* tool enables you to maintain details of the SAML identity provider. The following table indicates which information can be maintained.

General SAML IDP Details

UI Element	Description	Example
<i>Name</i>	The name of the SAML identity provider is mandatory and must be unique.	ACCOUNTS_ACME_COM
<i>Subject</i>	SAML IDP is specified in an XML document containing the IDP metadata	CN=CPS Production, OU=WebKm, O=ACME, L=Accra, C=GH
<i>Issuer</i>	SAML IDP is specified in an XML document containing the IDP metadata	CN=CPS Production, OU=WebKm, O=ACME, L=Accra, C=GH
<i>Entity ID</i>	The entity ID of the remote SAML party	accounts.acme.com
<i>Dynamic User Creation</i>	Enable or disable the dynamic creation of new SAML users.	Disabled

Destination

The *Destination* screen area in the *SAML Identity Provider* tool enables you to maintain details of the HTTP destination for the system hosting the SAML identity provider service. You must provide a base URL for the SAML IDP as well as further, more detailed, information about the location of the resources that provide the sign-on and sign-off services. The following table indicates which information can be maintained.

Details of the SAML IDP's HTTP Destination

UI Element	Description	Example
<i>Base URL</i>	The resource location where the SAML identity provider is reachable.	https://accounts.acme.com:443
<i>SingleSignOn URL (RedirectBinding)</i>	URL of the IDP endpoint for SSO requests using SAML redirect binding	/saml2/idp/sso/accounts.acme.com
<i>SingleSignOn URL (PostBinding)</i>	URL of the IDP endpoint for SSO requests using SAML post binding	/saml2/idp/sso/accounts.acme.com
<i>SingleLogout URL (RedirectBinding)</i>	URL of the IDP endpoint for single logout (SLO) requests using SAML redirect binding	/saml2/idp/slo/accounts.sap.com
<i>SingleLogout URL (PostBinding)</i>	URL of the IDP endpoint for single logout (SLO) requests using SAML post binding	/saml2/idp/slo/accounts.sap.com

SAML bindings describe a protocol used to transport SAML messages: both the requests and the responses. The following bindings are relevant for the configuration of the HTTP destination for the SAML identity provider.:

- Redirect binding
The SAML message is in the URL itself as a query parameter. Redirect bindings enforce limitations on the message and ZLIB compression is required.
- Post binding
The SAML message is transported inside an HTTP body in the `POST` parameter. There is no limitation on the message and no compression needed.

12.1.6.3 Modify an Existing SAML Identity Provider

SAP HANA supports the use of SSO authentication based on Security Assertion Markup Language (SAML) certificates. An identity provider is used by the service provider to authenticate the users signing in by means of SSO.

Prerequisites

- SAP HANA uses roles to determine the level of access to the features provided by the SAP HANA XS Administration Tool. To access the tools required to add an SAML identity provider, you must have a role based on the role template `sap.hana.xs.admin.roles::SAMLAdministrator`.
- You have access to the XML document containing the IDP metadata that describes the SAML identity provider (SAML IDP) you want to modify.

Context

To edit the details of an SAML identity provider, you can use the [SAP HANA XS Administration Tool](#), as described in the following steps:

⚠ Caution

To maintain the details of SAML identity providers, you must be familiar with the technical background of SAML SSO mechanisms and requirements.

Procedure

1. Start the [SAP HANA XS Administration Tool](#).

The [SAP HANA XS Administration Tool](#) tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must have the privileges required to perform SAML administration tasks.

2. Start the [SAML Identity Provider](#) tool.

Choose [SAML Identity Provider](#) in the list of tools displayed on the left-hand side of the [SAP HANA XS Administration Tool](#) window.

3. Select the SAML identity provider, whose details you want to modify.

The list of available SAML IDPs is displayed in the [SAML Identity Provider List](#) on the left-hand side of the [SAML Identity Provider](#) tool.

4. Modify the details of the selected SAML SSO identity provider (IDP).

The information required to maintain details of an SAML IDP is specified in an XML document containing the IDP metadata. This document should be available as part of the SAML service you want SAP HANA XS to use.

- a. Paste the contents of the XML document containing the IDP metadata into the [Metadata](#) box in the [Add Identity Provider Info](#) screen area.

If the contents of the XML document are valid, the parsing process extracts the information required to insert into the [Subject](#), [Entity ID](#), and [Issuer](#) fields in the [General Data](#) screen area, and the URL fields in the [Destination](#) screen area.

5. Save the modifications to the SAML identity provider.

Choose [Save](#) to save the changes.

Related Information

[Add an SAML Identity Provider \[page 1563\]](#)

12.1.7 Maintaining SMTP Server Configurations

Define details of the SMTP server that SAP HANA XS can use to respond to requests from applications to send e-mails.

The SMTP configuration defines the details of the SMTP server that is available for use by all applications running on an SAP HANA XS server. You can configure one SMTP server per SAP HANA XS server. As part of the configuration, you also specify the following options:

- General SMTP details
- Logon authentication type
- Transport-channel security type
- Other settings

SMTP Host System Details

When defining the details of the SMTP server to be used by the SAP HANA XS applications, you must specify the following elements:

- *Mail Server Host*
The name or the IP address of the system hosting the SMTP relay server that the XS applications can use to send an e-mail. The default value is *localhost*.
- *Mail Server Port*
The port number to use for connections to the SMTP relay server. The default value is *25*.

i Note

The port number to use can change according to the security type specified for the SMTP transport channel, for example, SSL or TLS.

SMTP Logon Authentication Type

You must tell SAP HANA XS which method the SMTP server uses to authenticate the logon credentials of the user that SAP HANA uses to establish the connection. The available choices for the authentication type are: *None*, *Auto*, *Logon*, *Plain*, *CRAM-MD5*, or *Digest-MD5*.

If you choose the option *None*, no logon credentials are required for the connection to the SMTP relay server. If you choose the option *Auto*, SAP HANA XS checks the authentication mechanisms supported by the SMTP relay server and selects one automatically according to the following order of preference: *Digest-MD5*, *CRAM-MD5*, *Plain*, *Login*, or *None*.

i Note

For all authentication-type options except *None*, you must specify the name and password of the user whose credentials SAP HANA XS uses to log on to the SMTP server.

SMTP Transport-Channel Security Type

When you set up the SMTP configuration, you must specify the security type used to encrypt the transport channel between the SAP HANA XS server and the SMTP server; you can choose any of the following values:

- *None*
This is default value for the transport security type; the channel used to communicate with the SMTP relay server is not encrypted. Note that it is possible that both SAP HANA XS and the specified SMTP relay server are running in the same trusted network or even on the same host.
- *STARTTLS*
You can specify STARTTLS as the transport security only if it is supported by the SMTP relay server. If it is not supported, the application trying to send an e-mail encounters an error and the requested e-mail message is not sent.
- *SSL/TLS*
Use an SSL/TLS-wrapped channel to communicate with the SMTP relay server. If SSL/TLS is not supported by the SMTP relay server then the connection cannot be established, the application trying to send an e-mail encounters an error, and the requested e-mail message is not sent. If you choose SSL/TLS as the transport security type, you will very probably have to specify a different port, usually 465, in the SMTP host section. You will also have to specify the name of the trust store holding the certificates and keys required to establish a trusted connection with the SMTP server.

Note

If the SMTP relay server's certificate cannot be verified, then the connection to the specified SMTP server cannot be established, the application trying to send an e-mail encounters an error, and the requested e-mail message is not sent.

Socket Proxy Settings

If your system uses a proxy service for Socket Secure (SOCKS) routing, you need to enable support using the SOCKS Proxy toggle button (ON) and, in addition, provide connection details for the system where the proxy service is running, for example:

Caution

The proxy-server settings you define here are overridden by any SAP HANA system wide setting for a proxy server, for example, defined by the `enforced_outbound_proxy` parameter in the `communication` section of the `xsengine.ini` configuration file.

- *Proxy Host*
The name of the system hosting the SOCKS proxy service
- *Proxy Port*
The port number to use for connections to the SOCKS proxy server running on the host specified in *Proxy Host*
- *Proxy Username*
The name of the user whose account is used to log on to the SOCKS proxy server system specified in *Proxy Host*

- [Proxy Password](#)

The password of the user whose account is used to log on to the SOCKS proxy server system specified in [Proxy Host](#)

Other Settings

You can specify the maximum length of time (in milliseconds) that SAP HANA XS must wait for a response from the SMTP relay server with which it is trying to establish a connection; the default value is 60000 milliseconds (1 minute). If the specified timeout limit is reached, the connection is reset and the application requesting the connection encounters an error.

i Note

If a connection is reset due to a timeout problem, the state of any sent e-mail messages is unknown. However, some useful information might be available in the logs of the SMTP relay server.

Related Information

[Create an SMTP Configuration \[page 1570\]](#)

[SAP HANA XS Classic Configuration Parameters \[page 1532\]](#)

12.1.7.1 Create an SMTP Configuration

Define the settings an SAP HANA XS application uses for outbound connections to an SMTP server.

Prerequisites

SAP HANA uses roles to determine the level of access to the features provided by the [SAP HANA XS Administration Tool](#). To access the [SMTP Configuration](#) tools that enable you to set up an SMTP server for SAP HANA XS applications, you must have roles based on the following role templates:

- `sap.hana.xs.admin.roles::RuntimeConfAdministrator`
- `sap.hana.xs.admin.roles::SMTPDestAdministrator`

Context

The SMTP configuration defines the details of the SMTP server that is available for use by all applications running on an SAP HANA XS server. You can configure one SMTP server per SAP HANA XS server. As part of

the configuration, you also specify what authentication type to use when establishing the connection as well as the security type used to encrypt the transport channel between the SAP HANA XS server and the SMTP sever, for example, SSL or TLS. To create an SMTP configuration for an SAP HANA XS application, perform the following steps:

Procedure

1. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must have the privileges required to perform administration tasks with the *SMTP Configurations* tool.

2. Start the *SMTP Configurations* tool.

In the list of XS Administration tools, choose *SMTP Configurations* to display the screen where you can manage the configuration of the SMTP server used by SAP HANA XS applications.

i Note

You can configure only one SMTP server for each SAP HANA XS instance.

3. Specify details of the system hosting the SMTP server that the SAP HANA XS applications must use.

Provide the name of the system hosting the SMTP server and the port number required to open a connection. The default value for the port number is 25.

i Note

If SSL or TLS is required to encrypt the transport channel, the port number will probably change, for example, to 465.

4. Specify the authentication settings required for access to the SMTP host.

Choose an authentication method from the *Authentication Type* drop-down list, for example, *auto*, *logon*, or *none* and, if necessary, the user credentials required to log on to the SMTP server.

→ Tip

If you choose *auto*, setup checks the authentication mechanisms supported by the specified SMTP server and selects one in the following order of preference: *Digest-MD5*, *CRAM-MD5*, *Plain*, *Login*, and *None*.

5. Specify the security settings for the transport-channel.

The transport channel is used for the communication between the SAP HANA XS application and the SMTP server. If you choose either the *STARTTLS* or the *SSL/TLS* option, use the *Trust Store* drop-down list to specify the trust store where the certificates and keys for the SMTP sever are located.

i Note

If you choose the option *None*, the channel used to communicate with the SMTP relay server is not encrypted.

6. Define the timeout setting for connections to the specified SMTP server.

You can specify the maximum length of time (in milliseconds) that SAP HANA XS must wait for a response from the SMTP server with which it is trying to establish a connection; the default value is 60000 milliseconds (1 minute).

i Note

If the specified timeout limit is reached, the connection is reset and the application requesting the connection encounters an error.

7. Define the socket proxy settings.

If your system uses a proxy service for Socket Secure (SOCKS) routing, you need to enable support using the *SOCKS Proxy* toggle button (*ON*) and, in addition, provide connection details for the system where the proxy service is running, for example, the host name, the port number to use for connections, and the user credentials required to log on.

⚠ Caution

The proxy-server settings you define here are overridden by any SAP HANA system wide setting for a proxy server, for example, defined by the `enforced_outbound_proxy` parameter in the `communication` section of the `xsengine.ini` configuration file.

8. Save the changes you have made to the SMTP configuration.

Related Information

[Maintaining SMTP Server Configurations \[page 1568\]](#)

[SMTP Configuration Details \[page 1572\]](#)

12.1.7.1.1 SMTP Configuration Details

The SMTP configuration defines the details of the SMTP server that is available for use by all applications running on an SAP HANA XS server.

As part of the SMTP configuration, you specify the following options:

- [General SMTP settings \[page 1573\]](#)
- [Logon authentication type \[page 1573\]](#)
- [Transport security type \[page 1573\]](#)
- [Socket proxy settings \[page 1574\]](#)
- [Other settings \[page 1574\]](#)

General SMTP settings

The *General SMTP Settings* screen area of the *SMTP Configurations* tool enables you to maintain the basic details of the system hosting the SMTP server that SAP HANA XS applications use to send e-mail. The following table indicates which information can be maintained.

UI Element	Description	Example
<i>Mail Server Host</i>	The name of the system hosting the SMTP server.	localhost
<i>Mail Server Port</i>	The port to connect to on the SMTP server; default is 25.	25

Authentication

The *Authentication* screen area of the *SMTP Configurations* tool enables you to maintain details of the user credentials required to log on to the system hosting the SMTP server and the mechanism used during the logon process to carry out user authentication. The following table indicates which information can be maintained.

UI Element	Description	Example
<i>Authentication Type</i>	The method used by the SMTP server to authenticate the credentials of the user that SAP HANA uses to establish the connection	None, Auto, Logon, Plain, CRAM-MD5, or Digest-MD5.
<i>Username</i>	For all authentication-type options except <i>None</i> , the name and password of the user whose credentials SAP HANA XS uses to log on to the SMTP server.	john doe
<i>Password</i>	For all authentication-type options except <i>None</i> , the password of the user whose credentials SAP HANA XS uses to log on to the SMTP server.	*****

Transport Security Settings

The *Transport Security Settings* screen area of the *SMTP Configurations* tool enables you to maintain details of the security type used to encrypt the transport channel between the SAP HANA XS server and the SMTP server. The following table indicates which information can be maintained.

UI Element	Description	Example
<i>Transport Security</i>	The method used by the SMTP server to authenticate the credentials of the user that SAP HANA uses to establish the connection	None, STARTTLS, SSL/TLS
<i>Trust Store</i>	Contains the certificates used to establish trust relationships between servers, for example, SAP HANA XS and the SMTP server	sapspv.pse

Socket Proxy Settings

The *Socket Proxy Settings* screen area of the *SMTP Configurations* tool enables you to maintain details of the system hosting the proxy service used by the SMTP server for Secure Socket (SOCKS) routing. The following table indicates which information can be maintained.

UI Element	Description	Example
<i>SOCKS Proxy</i>	Enable/Disable Socket Secure (SOCKS) routing	N/A
<i>Proxy Host</i>	Name of the system hosting the proxy service for Socket Secure (SOCKS) routing	smtp.host.acme.com
<i>Proxy Port</i>	Port number to use for connections to the proxy server	1080
<i>Proxy Username</i>	Name of the user required to log on to the proxy server	johndoe
<i>Proxy Password</i>	Password for the user required to log on to the proxy server	****

Other Settings

The *Other Settings* screen area of the *SMTP Configurations* tool enables you to maintain additional details of the SMTP server, for example, the connection timeout setting. The following table indicates which information can be maintained.

UI Element	Description	Example
<i>Timeout</i>	Maximum length of time (in milliseconds) that SAP HANA XS must wait for a response from the SMTP server with which it is trying to establish a connection	60,000 milliseconds (1 minute)

Related Information

[Create an SMTP Configuration \[page 1570\]](#)

12.1.8 Maintaining HTTP Access to SAP HANA

Ensure that Web-based applications have access to SAP HANA via HTTP.

To enable access to the services provided by the XS-based applications that you develop for SAP HANA, you need to ensure that client applications can access the SAP HANA XS Web server by HTTP or HTTPS. As part of the configuration process, you also need to configure SSL (for use with secure HTTP), set up the SAP Web Dispatcher (for example, to use non-default ports or secure HTTP), and maintain the trust stores that store the certificates required for secure communication. In addition, in a multi-database environment, you also need to configure HTTP access to multi-tenant database containers.

Maintaining HTTP access to SAP HANA includes one of more of the following tasks:

- Configure HTTPS (SSL) for client application access
Configure the SAP Web Dispatcher to use HTTPS (SSL) for incoming requests from UI front ends and applications, for example, SAP HANA applications. The requests are then forwarded by the SAP Web Dispatcher to SAP HANA.
- Maintain standard HTTP port numbers for SAP HANA XS
Check or change the default HTTP port settings, for example, to ensure that standard ports 80 and 443 are used for client access to the SAP HANA XS Web server by HTTP or HTTPS, respectively.
- Configure HTTP access to multi-tenant database containers
Configure the internal SAP Web Dispatcher so that, in an environment where multiple tenant database containers are available, the SAP Web Dispatcher knows which client requests to dispatch to which tenant database, for example, on the basis of alias DNS names.

Related Information

[Configure HTTPS \(SSL\) for Client Application Access \[page 1575\]](#)

[Maintain Standard HTTP Port Numbers with SAP HANA XS \[page 1577\]](#)

[Configure HTTP\(S\) Access to Tenant Databases via SAP HANA XS Classic \[page 1578\]](#)

12.1.8.1 Configure HTTPS (SSL) for Client Application Access

To improve the security of your SAP HANA landscape, you can configure the SAP Web Dispatcher to use HTTPS (SSL) for incoming requests from UI front ends and applications, for example, SAP HANA applications. The requests are then forwarded to SAP HANA.

Prerequisites

If you want to set up a secure SSL connection (Secure Socket Layer) between client applications and the SAP Web Dispatcher, the following components are prerequisites:

- The CommonCryptoLib library (`libsapcrypto.so`)
CommonCryptoLib (`libsapcrypto.so`) is installed by default as part of SAP HANA server installation at `$DIR_EXECUTABLE`.
- You have a role based on the role template `sap.hana.xs.wdisp.admin::WebDispatcherAdmin`. This is required to access the SAP HANA *Web Dispatcher Administration* tool.

Context

The SAP Web dispatcher lies between the Internet and your SAP system. It is the entry point for HTTP(s) requests into your system. To configure the SAP Web Dispatcher to use SSL for inbound application requests, perform the following steps.

Procedure

1. Start the *SAP HANA Web Dispatcher Administration* tool.

The *SAP HANA Web Dispatcher Administration* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/wdisp/admin/`.

Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must have the privileges required to perform administration tasks with the *Web Administration Interface* of the *SAP HANA Web Dispatcher Administration* tool.

2. Create an SSL key pair and a certificate request:

The SSL key pair is created with the default `SAPSSLS.pse` trust store; if you want to create a new SSL key pair, choose *Recreate PSE* in the *PSE Management* tool. To create a certificate request, perform the following steps:

- a. Open the *PSE Management* tool.

In the *SAP HANA Web Dispatcher Administration* tool, choose **SSL and Trust Configuration** **PSE Management**.

- b. Create the certificate request.

In the *PSE Management* screen area, choose *Create CA Request*.

- c. Submit the generated certificate request to your certificate authority (CA) for signing.

Copy the contents of the certificate request from the *CA Request of PSE SAPSSLS.pse* screen area and send it to your certificate signing authority.

3. Import the signed certificate.

Add a copy of the signed certificate to the `SAPSSLS.pse` trust store. The certificate-request response must be generated in the correct format, for example, PKCS#7 certificate chain, which contains both the requester's and the issuing CA's certificates. If the response contains only the requester's certificate in PEM (Privacy Enhanced Mail) format and no CA certificate, the system can still build the correct format. However, in this case, the issuing CA's root certificate must already be available in the same certificate store into which you import the requester's certificate.

Tip

Make sure that the date and time settings on the server hosting the SAP Web Dispatcher are correct and synchronized with the certificate authority (CA) that issued the certificate you import, otherwise the certificate might be interpreted as invalid.

- a. Open the *PSE Management* tool.

In the *SAP HANA Web Dispatcher Administration* tool, choose **SSL and Trust Configuration** > *PSE Management* .

- b. Select the target trust store.

In the *Manage PSE* screen area, choose *SAPSSLS.pse* from the drop-down menu.

- c. Import the signed certificate request.

In the *PSE Attributes* screen area, choose *Import CA Response* and copy the signed certificate response from your CA into the *Import CA Request of PSE SAPSSLS.pse* screen area.

12.1.8.2 Maintain Standard HTTP Port Numbers with SAP HANA XS

The default HTTP port settings for SAP HANA XS include an SAP HANA instance number, for example, 80<*SAP HANA instance*> (8000). You can change the default settings, for example, to ensure that standard ports 80 and 443 are used for client access to the SAP HANA XS Web server by HTTP or HTTPS.

Prerequisites

To configure the SAP HANA XS server to use the standard HTTP ports 80 and 443, bear in mind the following prerequisites:

- Superuser authorization is required to bind ports with a number less than (<) 1024 (well-known ports) on a UNIX system
- Neither the ICM process nor the SAP Web Dispatcher has the superuser authorization.

Context

By default, the SAP HANA XS Web server is configured to use the port numbers 80<*SAP HANA instance number*> for HTTP and 43<*SAP HANA instance number*> for HTTPS requests from clients. You can change this behavior, for example, to configure the SAP HANA XS server to use the standard HTTP ports 80 and 443, as follows:

Procedure

1. Open the instance profile of your SAP Web Dispatcher.

The SAP Web Dispatcher profile can be found in the following location in the SAP HANA studio:

SSL and Trust Configuration > *Configuration* > *webdispatcher* > *[profile]* 

2. Check and, if necessary, modify the HTTP/S port settings in the SAP Web Dispatcher profile, as follows:

```
icm/server_port_0 = PROT=HTTP, PORT=80, PROCTIMEOUT=600, EXTBIND=1
icm/server_port_1 = PROT=HTTPS, PORT=443, PROCTIMEOUT=600, EXTBIND=1
```

Save the changes to the SAP Web Dispatcher profile.

3. Bind the default SSL port to use.

Since only users with superuser authorization rights can bind ports with a number less than (<) 1024 (well-known ports) on a UNIX system, and the ICM process or the SAP Web Dispatcher should not have these rights (and ICM cannot have them for technical reasons), the port must be bound by an external program and the listen socket then transferred to the calling process. You can use the `icmbnd` command.

i Note

The installation process creates the file `icmbnd.new`, which you must rename to `icmbnd` and configure as described below. This applies after a system update, too.

Since superuser privileges are required to bind ports with a number lower than 1024, you must change the owner and permissions of the `icmbnd` command, for example, from `<SID>adm` to user `root`.

- a. Change the owner of the `icmbnd` command:

```
$> chown root:sapsys icmbnd
```

- b. Change the permissions for the `icmbnd` command:

```
$> chmod 4750 icmbnd
```

- c. Check the new permissions for the `icmbnd` command:

```
$> ls -al
rwsr-x 1 root sapsys 1048044 Feb 13 16:19 icmbnd
```

Related Information

[SAP Web Dispatcher: Binding Ports < 1024 on UNIX](#)

12.1.8.3 Configure HTTP(S) Access to Tenant Databases via SAP HANA XS Classic

To enable Web-based applications to send HTTP(S) requests to tenant databases via the SAP HANA XS classic server, the internal SAP Web Dispatcher must be configured so it knows which requests to dispatch to which database on the basis of DNS alias host names. You do this by specifying the public URL of every tenant database in the `xsengine.ini` configuration file.

Prerequisites

- You are logged on to the system database.
- You have the system privilege INIFILE ADMIN.
- The network administrator has defined an alias hostname in your organization's Domain Name System (DNS) for every tenant database in the SAP HANA system. The alias hostname must refer to the hostname of the machine that is used for HTTP(S) access to the tenant database.
- You have a role based on the role template `sap.hana.xs.wdisp.admin::WebDispatcherAdmin`. This is required to access the SAP HANA Web Dispatcher Administration tool for configuring HTTPS.

Context

The XS classic server allows Web-based applications to access SAP HANA via HTTP(S). The internal Web Dispatcher of the SAP HANA system manages these incoming HTTP(S) requests. To allow applications to send requests to specific databases, every tenant database needs an alias host name. Requests to the alias host name can then be forwarded to the XS server of the corresponding tenant database. Requests with the physical host name in the HTTP host header are forwarded to the XS server running on the system database.

The default HTTP ports are used in all cases, that is, 80<instance> (HTTP) and 43<instance> (HTTPS). Alias host names are mapped to internal HTTP(S) ports so that incoming requests can be routed to the correct database.

You configure HTTP(S) access to tenant databases by specifying in the `xsengine.ini` file the URLs by which each tenant database is publicly accessible. The system then automatically configures the Web Dispatcher by generating the required profile entries in the `webdispatcher.ini` configuration file. It is not necessary to specify the URL of the system database, this is done automatically.

Note

This automatic configuration of the Web Dispatcher is controlled by the parameter `[profile] wdisp/system_auto_configuration` in the `webdispatcher.ini` configuration file. If this parameter is set to **false**, you need to configure the `webdispatcher.ini` file manually.

For HTTPS access, you must subsequently configure the required client certificates and trust stores using the SAP Web Dispatcher Administration tool. The following approaches are supported:

- Using a single "wildcard" server certificate in a single trust store that covers all databases in the system. Wildcard certificates are more flexible when tenant databases are frequently added and deleted. However, if you use a wildcard certificate, either the server requires its own sub-domain or you must ensure that the certificate cannot be abused from other servers.

Caution

Do not use a wildcard server certificate if strict isolation between tenant databases is required. If authentication relies on a wildcard certificate and a shared trust store, users of one tenant database will be able to log on to other databases in the system.

- Using individual certificates in individual trust stores for each database

Individual certificates for each database are more suitable in a flat domain structure for individual servers. They also ensure strict isolation between tenant databases. However, they involve more administrative effort to maintain.

Procedure

1. Specify the public URLs of all tenant databases in the `xsengine.ini` file in one of the following ways:

Option	Description
SAP HANA studio	<ol style="list-style-type: none"> 1. Open the Administration editor and choose the <i>Configuration</i> tab. 2. Navigate to the <code>xsengine.ini</code> file and expand the <code>public_urls</code> section. 3. For each tenant database in the system, add the new properties <code>http_url</code> and <code>https_url</code> at the database layer and enter its public URL as the value: <code>http://<virtual_hostname>:80<instance></code>
SQL	<p>For each tenant database, execute the statements:</p> <ul style="list-style-type: none"> ◦ ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', '<tenant_DB_name>') SET ('public_urls', 'http_url') = 'http://<virtual_hostname>:80<instance>' WITH RECONFIGURE; ◦ ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', '<tenant_DB_name>') SET ('public_urls', 'https_url') = 'https://<virtual_hostname>:43<instance>' WITH RECONFIGURE;

Note

The following values are set at the **default layer** and represent the URLs of the system database:

- `http://$ (SAPLOCALHOST) : 80$ (SAPSYSTEM)`
- `https://$ (SAPLOCALHOST) : 43$ (SAPSYSTEM)`

By default, the system database initially retrieves any request with the port `80<instance_no>`. However, as soon as you configure the URLs of tenant databases, it is available under `http://<localhost>:80<instance>` only, and not the fully qualified domain name (FQDN). The local host is known to SAP HANA without the FQDN.

If you want to change this default behavior and configure a different URL for the system database, you can do so by executing the following statement:

```
ALTER SYSTEM ALTER CONFIGURATION ('nameserver.ini', 'system')
SET('public_urls', 'http_url') = 'http://<virtual_hostname>:80<instance>'
WITH RECONFIGURE;
```

New entries are now created in the `webdispatcher.ini` file at the host layer for every database. You can verify this by executing the following statement (from the system database):

```
SELECT KEY, VALUE, LAYER_NAME FROM SYS.M_INIFILE_CONTENTS WHERE FILE_NAME =
'webdispatcher.ini' AND SECTION = 'profile' AND KEY LIKE 'wdisp/system%'
```

This returns the following result for example:

```
KEY          | VALUE                                                                 | LAYER_NAME
wdisp/system_0 | GENERATED, SID=SYS, EXTSRV=http://localhost:30014, SRCVHOST='myhost' | DEFAULT
wdisp/system_1 | GENERATED, SID=MYD, EXTSRV=http://localhost:30042, SRCVHOST='mydatabase.example.com' | HOST
```

- Optional: Secure incoming communication by configuring HTTPS.

Option	Description
Single certificate for all databases	<ol style="list-style-type: none"> Start the SAP HANA Web Dispatcher Administration tool at <code>http://<localhost>:80<instance>/sap/hana/xs/wdisp/admin/</code>. For the default <code>SAPSSLS.pse</code> trust store, create a new SSL key pair and certificate request: <ol style="list-style-type: none"> From the main menu, choose SSL and Trust Configuration > PSE Management. From the Manage PSE menu, choose SAPSSLS.pse. Choose Recreate PSE. Enter a distinguished name that matches the host name of all tenant databases. <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p>❖ Example</p> <ul style="list-style-type: none"> ○ Physical host name: myhost.example.com ○ Tenant host name 1: mydatabase1.example.com ○ Tenant host name 2: mydatabase2.example.com <p>In this case, you specify CN=*.example.com as the DN, thus creating a server certificate that matches all tenant databases and the system database.</p> </div> <ol style="list-style-type: none"> Choose Create. Create a certificate request and submit to your certificate authority (CA) for signing (Create CA Response). Import the signed certificate <p>For more information, see <i>Configure HTTPS (SSL) for Client Application Access</i>.</p>

Individual certificates for each database	<ol style="list-style-type: none"> Start the SAP HANA Web Dispatcher Administration tool at <code>http://<localhost>:80<instance>/sap/hana/xs/wdisp/admin/</code>. For each tenant database and the system database, create a new trust store with a unique certificate: <ol style="list-style-type: none"> From the main menu, choose SSL and Trust Configuration > PSE Management. On the PSE management screen, choose Create New PSE. Enter a file name for the new PSE. <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <p>❖ Example</p> <p>example.pse</p> </div> <ol style="list-style-type: none"> Enter the distinguished name: <p>CN=<host name used for the tenant database in the public_urls section of the xsengine.ini file></p> Choose Create. For the new PSE, create a certificate request and submit to your CA for signing (Create CA Response). Import the signed certificate into the new PSE (Import CA Response). Configure the Web Dispatcher to use multiple certificates: <ol style="list-style-type: none"> In the <code>webdispatcher.ini</code> file, create or change the parameter <code>[profile] icm/ssl_config_0</code>, specifying as the value: <p>ID=ssl_config_main, CRED=SAPSSLS.pse, SNI_CREDS=<semicolon (';') separated list of database PSE files></p> Add ,SSLCONFIG=ssl_config_main to the value of the <code>icm/server_port</code> parameter for the HTTPS port (by default <code>icm/server_port_1</code>).
--	---

Option	Description
	<p>❖ Example</p> <pre>icm/server_port_1 = PROT=HTTPS,PORT=4443\$(SAPSYSTEM),PROCTIMEOUT=600, SSLCONFIG=ssl_config_main</pre>

Results

You can access the XS server of tenant databases via the configured URLs.

→ Tip

If you experience slow response times when accessing the XS server of a tenant database (for example, Web-based applications running on the tenant database), this indicates that the server is not able to resolve the host name correctly using the DNS and retries repeatedly. If this is the case, contact your network administrator for a detailed problem analysis.

As a workaround, you can manually override virtual host name resolution on the machine where the browser is running by modifying the `/etc/hosts` file on the local machine. In this file, append a new line, starting with the static IP address of the server, followed by the virtual host name of your tenant database, for example, "10.20.30.40 mydatabase.example.com". To edit this file you need admin or root privileges.

Next Steps

Optional: Enable access to Web-based applications from the SAP HANA studio.

Some Web-based tools are accessible from the SAP HANA studio, for example, the SAP HANA cockpit and SAP HANA Lifecycle Management tool. If you want to be able to access these tools from a tenant database registered in the studio, you must specify the alias hostname in the properties. You can do this as follows:

1. In the *Systems* view, right-click the tenant database and choose *Properties*.
2. Open the *XS Properties* page and enter the alias hostname in the *XS Host* field.

Related Information

[Configure HTTPS \(SSL\) for Client Application Access \[page 1575\]](#)

[Using SAP Web Dispatcher for Load Balancing with Tenant Databases \[page 280\]](#)

12.1.9 Maintaining Single Sign-On for SAP HANA XS Applications

You can configure SAP HANA applications to use single sign-on (SSO) authentication to confirm the logon credentials of a user calling an application service. SAP HANA supports SSO certificates based on the Security Assertion Markup Language (SAML) or X.509.

If you want your the SAP HANA XS applications to use an SSO certificate based on SAML or X.509 as the logon authentication method, you must perform the following high-level steps:

1. Maintain the SAP HANA trust store.

SAP HANA makes use of multiple trust stores; the trust stores listed below are provided by default.

i Note

The trust stores listed below are located in the **file system**. In some cases, it is possible and recommended to use trust stores that exist in the database as database objects. In-database trust stores (referred to as certificate collections) contain the required client certificates, which are also stored in the database. We recommend using in-database certificate collections where possible. For more information, see *Managing Client Certificates in the SAP HANA Database*.

- The SAP HANA trust store (`sapsrv.pse`)
Used for secure SQL and SAML or OAuth scenario, `sapsrv.pse` is installed automatically and is available by default.

→ Recommendation

For user authentication based on X.509 certificates and SAML assertions, we recommend creating separate certificate collections with the purposes [X.509](#) and [SAML](#) instead of using the file system-based trust store `sapsrv.pse`.

- The SAP Web Dispatcher trust store (`SAPSSLS.pse`)
Required for secure connections using the Secure Socket Layer (SSL) protocol, `SAPSSLS.pse` is installed automatically and is available by default.
- The SAP Logon Ticket trust store (`saplogon.pse`)
Optional: `saplogon.pse` is only necessary if an SAP HANA XS application requires an SAP logon ticket from a user at logon.

→ Recommendation

For user authentication based on logon tickets, we recommend creating a certificate collection with the purpose [SAP LOGON](#) instead of using the file system-based trust store `saplogon.pse`.

- The client authentication trust store (`SAPSSLC.pse`)
Optional: `SAPSSLC.pse` is only required for client connections, for example, that use the SQL client interface (`hdbsql`).
2. Choose the SSO authentication method and configure the trust relationships:
Trust relationships are required between SAP HANA and any remote system providing services that an SAP HANA XS application requires.
 - SSO with X.509 certificates
Add the root certificate of the Certification Authority (CA) for trusted X.509 certificates to both the SAP HANA trust store **and** the trust store for the SAP Web Dispatcher.

- SSO with SAML certificates
Obtain, authenticate, and import the SAML identity provider (IDP) metadata (an XML document) for the SAML service provider.
- 3. Maintain the SSO provider for SAP HANA XS
Maintain a runtime configuration for the SAP HANA application, which indicates that user authentication is by means of SSO certificates based on either SAML or X.509.

Related Information

[Managing Client Certificates \[page 900\]](#)

[Configure SSO with X.509 Authentication for SAP HANA XS Applications \[page 1584\]](#)

[Configure SSO with SAML Authentication for SAP HANA XS Applications \[page 1586\]](#)

12.1.9.1 Configure SSO with X.509 Authentication for SAP HANA XS Applications

SAP HANA applications can use single sign-on (SSO) authentication with X.509 certificates to confirm the logon credentials of a user calling an application service.

Prerequisites

- You have a role based on the role template `sap.hana.xs.admin.roles::RuntimeConfAdministrator`.
- The CommonCryptoLib library (`libsapcrypto.so`) is installed and available.
- A certificate collection with the purpose *X.509* is available. For more information, see *Managing Client Certificates in the SAP HANA Database*.
- The SAP Web Dispatcher trust store (`SAPSSLS.pse`) is available.

Context

To enable SAP HANA applications to use single sign-on (SSO) authentication with X.509 certificates to confirm the logon credentials of a user, you need to add the root certificate of the Certification Authority that issues trusted X.509 certificates to both the SAP HANA trust store for X.509 authentication and the trust store of the SAP Web Dispatcher, `SAPSSLS.pse`.

Procedure

1. Add the root certificate (for example, `SSO_CA.der`) to the SAP HANA trust store, that is the certificate collection with purpose *X.509*.
 - a. Open the SAP HANA cockpit.
 - b. Open the *Certificate Store* app.
 - c. Import the root certificate into the certificate store.
 - d. Open the *Certificate Collections* app.
 - e. Select the collection with purpose *X.509*.
 - f. Add the root certificate to this collection.
2. Add the root certificate (for example, `SSO_CA.der`) to the SAP Web Dispatcher trust store (`SAPSSLS.pse`).
 - a. Start the *SAP HANA Web Dispatcher Administration* tool.
 - b. Open the *PSE Management* tool.

In the *SAP HANA Web Dispatcher Administration* tool, choose **SSL and Trust Configuration** > *PSE Management* .
 - c. Specify the trust store (PSE file) for the import operation.

In the *PSE Management* screen area, choose `SAPSSLS.pse` from the *Manage PSE* drop-down list.
 - d. Import the `SSO_CA.der` certificate.

In the *Trusted Certificates* screen area, choose *Import Certificate*.

Alternatively, you can also use the `sapgenpse` tool to import the `SSO_CA.der` certificate.

```
./sapgenpse maintain_pk -p /usr/sap/<SAPHANAInstance>/HDB<InstNo>/  
<Hostname>/sec/SAPSSLS.pse -a SSO_CA.der
```

3. Maintain the authentication settings in the runtime configuration for your SAP HANA XS application. You can use the Web-based SAP HANA XS Administration *Trust Manager* tool to complete this step. The tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

Note

The user maintaining the security settings needs the privileges granted by the SAP HANA XS role *TrustStoreAdministrator*.

4. Create a database user whose identity is defined in an X.509 certificate issued by your CA.
 - a. Create a new user in the SAP HANA database based on the details specified in an existing X.509 certificate.

The following example shows how to use the SQL statement `CREATE USER WITH IDENTITY` to create the database user "MyUserName" and the corresponding X.509 certificate:

```
CREATE USER MyUserName WITH IDENTITY 'CN=MyUserName, O=SAP-AG, C=DE' ISSUER  
'CN=SSO_CA, O=SAP-AG, C=DE' FOR X509
```
 - b. Import into the Web browser the X.509 certificate that is to be used to authenticate the new database user.
5. Use a Web browser to test the logon authentication settings for the SAP HANA application. When you enter the URL for your application in the Web browser, the Web browser prompts you to select a certificate, which enables you to log on without supplying logon credentials manually.

Related Information

[Managing Client Certificates \[page 900\]](#)

12.1.9.2 Configure SSO with SAML Authentication for SAP HANA XS Applications

SAP HANA applications can use single sign-on (SSO) authentication with SAML assertions to confirm the logon credentials of a user calling an application service. SAML assertions are certificates that comply with the Security Assertion Markup Language.

Prerequisites

- You have an advanced understanding of how SAML works.
- The CommonCryptoLib library (`libsapcrypto.so`) is installed and available on the SAP HANA server.
- You are authorized to edit the certificate collection with purpose [SAML](#) exists. You need system privilege `CERTIFICATE ADMIN` and object privilege `ALTER` on the collection. For more information, see *Managing Client Certificates in the SAP HANA Database*
- An SAML identity provider (IDP) is available and the corresponding SAML metadata (in the form of an XML document). For more information see *Add an SAML Identity Provider*.
- You have root/administrator access to the SAP HANA system that is configured to act as an SAML **service** provider.
- To maintain security and authentication settings for SAP HANA XS applications, you must have a role based on the role template `sap.hana.xs.admin.roles::RuntimeConfAdministrator`. To maintain SAML settings for SAP HANA XS applications, you need a role based on the role template `sap.hana.xs.admin.roles::SAMLAdministrator`.

Context

To enable SAP HANA applications to use single sign-on (SSO) authentication with SAML assertions to confirm the logon credentials of a user, you must copy the SAML certificate from the SAML IDP metadata document and add the certificate to the SAP HANA trust store for SAML authentication.

Procedure

1. Gather the metadata for the SAML identity provider (IDP).
This SAML IDP metadata typically takes the form of an XML document, which you can obtain from your security system administrator.

2. Extract the certificate string (which is DER encoded) from the SAML IDP metadata document. The certificate string is located in the `ds:X509Certificate` tag. For the SAP ID service, the certificate string could look like the following (incomplete) code example:

```
MIICHTCCAYagAwIBAgIETKTcJjANBgkqhkiG9w0BAQUFADBTMQswCQYDVQQGEwJERTEPMA0G...
```

3. Paste the extracted SAML certificate string into a file called `sapid.cer`.
4. Add the BEGIN and END tags to the SAML certificate.

The following example of a SAML certificate is incomplete; it is intended for illustration purposes only.

```
-----BEGIN CERTIFICATE-----  
MIICHTCCAYagAwIBAgIETKTcJjANBgkqhkiG9w0BAQUFADBTMQswCQYDVQQGEwJERTEPMA0G...  
-----END CERTIFICATE-----
```

5. Import the contents of the SAML certificate (`sapid.cer`) into the SAP HANA trust store, that is the certificate collection with purpose [SAML](#).
 - a. Open the SAP HANA cockpit.
 - b. Open the [Certificate Store](#) app.
 - c. Import the SAML certificate (`sapid.cer`) into the certificate store.
 - d. Open the [Certificate Collections](#) app.
 - e. Select the collection with purpose [SAML](#).
 - f. Add the SAML certificate (`sapid.cer`) to this collection.

6. Configure your SAP HANA system to act as an SAML service provider. For more information, see [Configure an SAP HANA System as an SAML Service Provider](#).

7. Maintain the authentication settings in the runtime configuration for the SAP HANA XS application for which you want to enable SSO with SAML authentication.

You can use the Web-based [SAP HANA XS Administration Tool](#) to complete this step. The tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>`:

`80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

The user maintaining the authentication settings in an application's runtime configuration needs the privileges granted by the SAP HANA XS role [RuntimeConfAdministrator](#).

Related Information

[Maintaining SAML Providers \[page 1559\]](#)

[Managing Client Certificates \[page 900\]](#)

12.1.9.3 Configure SSO with SAP Logon Tickets for SAP HANA XS Applications

SAP HANA applications can use single sign-on (SSO) authentication with SAP logon tickets to confirm the logon credentials of the user calling an application service.

Prerequisites

- You have administrator access to the SAP HANA system hosting the applications to which you want to enable access with SAP logon tickets.
- To maintain security and authentication settings for SAP HANA XS applications, you must have a role based on the role template `sap.hana.xs.admin.roles::RuntimeConfAdministrator`.
- The CommonCryptoLib library `libsapcrypto.so` is installed and available.
- A certificate collection with the purpose *SAP LOGON* is available. For more information, see *Managing Client Certificates in the SAP HANA Database*.

Context

To enable SAP HANA applications to use single sign-on (SSO) authentication with SAP logon tickets to confirm the logon credentials of a user requesting an application service, you must ensure that an SAP server is available that can issue SAP logon tickets. In addition, you need to add the server certificate of the ticket-issuing system to the SAP HANA trust store for authentication using logon tickets.

Procedure

1. Add the server certificate of the SAP system that issues SAP logon tickets to the SAP HANA trust store, that is the certificate collection with purpose *SAP LOGON*.
 - a. Open the SAP HANA cockpit.
 - b. Open the *Certificate Store* app.
 - c. Import the server certificate of the ticket-issuing system into the certificate store.
 - d. Open the *Certificate Collections* app.
 - e. Select the collection with purpose *SAP LOGON*.
 - f. Add the server certificate of the ticket-issuing system to this collection.
2. In SAP HANA, configure the details of the server that issues SAP logon tickets.

This step is optional but ensures that an SAP logon ticket can always be obtained in those cases where no SAP logon ticket is immediately available for the user trying to log on.

xsengine.ini		◆
application_container		
authentication		◆
logonticket_redirect_url		● https://vmw.sap.com:44333/sap/bc/

- a. Start the SAP HANA studio and open the *Administration* perspective.
- b. In the *Configuration* tab, expand (or add) the section `xsengine.ini` *authentication*.
- c. Set (or add) the parameter: `logonticket_redirect_url`.

Enter the URL that points to the system and service issuing SAP logon tickets, for example:

```
https://<hostname>:<portnumber>/<path/to/logon_ticket/service>
```

- o `<hostname>`
The hostname of the server issuing/storing the SAP logon tickets
- o `<portnumber>`
The port number accepting connections on the target server issuing/storing the SAP logon tickets
- o `</path/to/logon_ticket/service>`
Path to the service on the target system which handles the request for the SAP logon ticket. You can write your own custom ABAP service to handle these requests.

For example, the following URL would enable access to the **custom** (user-defined) SAP logon ticket service `zredirectwlogon` using port 44333 on the ABAP server `host.acme.com`:

```
https://host.acme.com:44333/sap/bc/zredirectwlogon?sap-client=<SAPClientNr>
```

3. Maintain the runtime configuration for the application that you want to use SAP logon tickets for user authentication.

You can use the Web-based *SAP HANA XS Administration Tool* to complete this step. The tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`. Choose *XS Artifact Administration*.

Note

The user maintaining the security settings needs the privileges granted by the SAP HANA XS role *RuntimeConfAdministrator*.

- a. Locate the root package of the application whose runtime configuration you want to modify.
Use the *Packages* list in the *Application Objects* pane.
- b. In the *Security & Authentication* tab, enable support for *SAP Logon/Assertion Ticket*.
- c. Save the changes you have made.

12.1.9.4 Configure Outbound SSO with Assertion Tickets

Assertion tickets are a form of bearer token that one application server uses to identify and authenticate a user on another application server, for example, in a single-sign-on (SSO) scenario. You can set up SAP HANA to function as the provider of the assertion tickets required to log on to a remote SAP server.

Prerequisites

To configure SAP HANA to use SAP assertion tickets to authenticate users who log on with SSO, note the following prerequisites:

- Your SAP HANA system is configured to use SSL
- You have administrator access to the SAP HANA system hosting the applications to which you want to enable access with SAP assertion tickets.
- To maintain security and authentication settings for SAP HANA XS applications, you must have a role based on the role template `sap.hana.xs.admin.roles::RuntimeConfAdministrator`. To maintain an HTTP destination, you need a role based on the role template `sap.hana.xs.admin.roles::HTTPDestAdministrator`.
- You know the system ID (SID) and client number of the SAP HANA system
- You know the system ID (SID) and client number of the remote SAP ABAP server that hosts the HTTP service (assertion-ticket provider) used by your XSJS application
- You have the permissions required to run transaction **STRUSTSSO2** in the ABAP system with which you want to establish a trust relationship.
- The CommonCryptoLib library `libsapcrypto.so` is installed and available on your SAP HANA system.
- You have read SAP Note [1982597](#)  concerning SAP logon tickets and assertion tickets which are created with UTF-8.

Context

SAP HANA XS enables you to build XSJS applications that use single sign-on services with authentication using SAP assertion tickets to consume additional Web services, for example, provided by a remote ABAP application server. If the XSJS application service requires access to remote services, you can create an HTTP destination that defines the logon details required by the remote ABAP system and specifies SSO with SAP assertion tickets as the logon authentication method. The assertion ticket is included in the header of the HTTP request sent by the application service; the remote system reads the HTTP header and uses the assertion to log the requesting user on automatically.

Procedure

1. Create the SAP HANA trust store for the assertion tickets, for example, `saplogonSign.pse`.

This trust store is used to issue the assertion tickets required for automatic logon to remote SAP systems using SSO.

```
sapgenpse gen_pse -p saplogonSign.pse "CN=<HOST>.<DOMAIN>, OU=<INSTANCE>, O=<ORG>, C=<COUNTRY>"
```

You are prompted to have the ticket signed by a Certificate Authority (CA):

- a. Copy the certificate request and have it signed by a known CA service.
- b. Copy the signed certificate results from the CA to the directory `/usr/sap/<SID>/HDB<Instance Number>/<machine name>/sec` on your SAP HANA system and name the file `saplogonSign.cer`.
- c. Import the signed certificate into the trust store.

```
./sapgenpse import_own_cert -c saplogonSign.cer -p saplogonSign.pse -r SAPNetCA.cer
```

2. Export the certificate that SAP HANA uses to sign assertion tickets.

You need to save the exported certificate to a local file for future use.

- a. Export the SAP HANA certificate from the SAP HANA trust store, for example, using the following command:

```
sapgenpse export_own_cert -p saplogonSign.pse
```

- b. Copy the output to a local file on your system.

3. Set up the trust relationship between SAP HANA and the remote SAP ABAP system you want to enable automatic logon with SSO and assertion tickets.

The remote SAP system hosting the HTTP service you want your XSJS application to use must trust the SAP HANA system hosting the XSJS service itself and acting as a provider of SAP assertion tickets.

- a. Log on to the target ABAP system and run transaction **STRUSTSSO2**.
- b. Select the system PSE (trust store).
- c. Choose the *import certificate* button in the certificate section.
- d. Select the SAP HANA certificate you signed in the previous step and import it.
- e. Choose the *Add to certificate list* button.
- f. Choose the *Add to ACL* button.
- g. Provide the system ID (SID) for the SAP HANA system; the client number is 000.
- h. Save the configuration.

4. Import the certificate of the system you want to trust for inbound SSO.

i Note

This step is optional; it is only required if you want to use SAP logon tickets for inbound SSO requests, too.

5. On the SAP HANA system, edit the configuration variable used to specify the name of the trust store for SAP assertion tickets.

Start the SAP HANA studio's *Administration Console* perspective and edit the parameter `saplogontickettruststore`. You can find the `saplogontickettruststore` parameter in

► `[indexserver | xsengine].ini` ► *authentication* ► `saplogontickettruststore` ►.

indexserver.ini		
[] authentication		
saml_service_provider_name		● http://localhost6.locald...
saplogontickettrace		● true
saplogontickettruststore		● saplogonSign.pse
session_cookie_validity_tir	180	

6. Maintain an HTTP destination for the XSJS service that needs access to a remote SAP system and set the authentication type to *SAP Assertion Ticket*.

You define the details of an HTTP destination in a configuration file that requires a specific syntax. The configuration file containing the details of the HTTP destination must have the file extension `.xshttpdest`.

⚠ Caution

The HTTP destination configuration and the XSJS application that uses it must reside in the same application package. An application cannot reference an HTTP destination configuration that is located in another application package.

- a. Create a plain-text file called `<MyHTTPdestination>.xshttpdest` and open it in a text editor.
- b. Use the following code to help you define the HTTP destination details.

i Note

Change the entries for the host name, port, system ID and client to suit your own requirements.

```
host = "<ABAP_server_name>";
port = <ABAP_HTTPS_PortNumber>;
description = "my SAP assertion ticket target";
useSSL = true;
pathPrefix = "";
authType = AssertionTicket;
useProxy = false;
proxyHost = "";
proxyPort = 0;
timeout = 0;
remoteSID = "<ABAP_SID>";
remoteClient = "<ABAP_ClientNumber>";
```

- c. Save and activate the file.

i Note

By default, saving the modified file automatically commits the saved version to the repository; you do not need to commit the file before activating it.

7. View the activated HTTP destination.
You can use the *SAP HANA XS Administration Tool* to check the contents of an HTTP destination configuration.

i Note

To make changes to the HTTP Destination configuration, you must use a text editor, save the changes and reactivate the file.

- a. Open a Web browser.
- b. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

To access details of HTTP destinations in the *SAP HANA XS Administration Tool*, you must have a role based on the role template `sap.hana.xs.admin.roles::HTTPDestAdministrator`.

- c. Locate the package containing the HTTP destination `<MyHTTPdestination>.xshttpdest`.
Expand the nodes in the *Application Objects* pane to locate the package where the HTTP destination resides and select the HTTP destination to display details in the right pane.
8. Check the specified system ID (SID) and the client of the remote SAP system referenced in the HTTP destination.
 - a. Enable the *SAP Assertion Ticket* radio button.
 - b. Check (or enter) the SID and client number for the remote SAP system in the *SAP SID* and *SAP Client* text boxes respectively.
9. Save the changes to the HTTP destination and use it in an XSJS application service.

→ Tip

You can reference an HTTP destination from an XSJS service using the function `$.net.http.readDestination("<packageName>", "<HTTPDestinationName>")`

12.1.10 Maintaining User Self Service Tools

User self-service tools enable SAP HANA users to trigger account-related tasks, for example, the creation of a new database account.

By default, the user self-service tools are disabled. The SAP HANA administrator must activate the user self-service feature to provide users with access to embedded tools they can use to request the creation of a new user account in the SAP HANA database or request a new password.

Setting up and maintaining user-self-service tools for SAP HANA includes the following high-level tasks:

- Enable user self-service tools
- Request a new user account
- Display a list of the current user requests
- Reject a user/user request
- Enable access to the user-self-service administration tool

The *USS Administration* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/selfService/admin`

i Note

To log on, use the name and password of the user who has a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

Enabling and maintaining the tools required to manage user self-service requests in SAP HANA involves the creation of a dedicated technical user and the assignment of dedicated roles.

- **Administrator**
The user who manages the self-service requests and access lists must be assigned a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`; the user self-service administrator is the same user as the user associated with the email address defined in the `xsengine.ini` parameter `sender_email`. The self-service administrator receives an e-mail in response to each self-service request; the e-mail contains a list of tasks to perform.
- **Technical user**
A dedicated technical user, who is used to execute tasks associated with user self-service requests, for example, sending e-mails in response to user requests. Technical users cannot be used to log on to SAP HANA.

Related Information

[Enable User Self-Service Tools \[page 1594\]](#)

[User Self-Service Roles \[page 1596\]](#)

12.1.10.1 Enable User Self-Service Tools

User self-service tools are not enabled by default; they must be activated by the SAP HANA administrator.

Prerequisites

To enable user self-service tools in SAP HANA, you must have the following privileges:

- Access to SAP HANA as SAP HANA database administrator
- Access to specific features provided by the SAP HANA XS administration tools, which requires the privileges granted by the following roles:
 - `sap.hana.xs.admin.roles::RuntimeConfAdministrator`
 - `sap.hana.xs.admin.roles::SQLCCAdministrator`
 - `sap.hana.xs.admin.roles::SMTPDestAdministrator`
 - `sap.hana.xs.ide.roles::SecurityAdmin`
 - `sap.hana.xs.selfService.admin.roles::USSAdministrator` (to log on to the user self-service administration tool)
- Access to the following SAP HANA tools:
 - [SAP HANA XS Administration Tool](#)
 - [SAP HANA Web-based Development Workbench](#)
 - [SAP HANA USS Administration Tool](#) (user self-service administration tool)

Context

By default, SAP HANA user self-service tools are disabled; the tools are neither visible in the user interface nor configured in SAP HANA. To provide access to embedded tools that enable users to request the creation of a new user account in the SAP HANA database or set a new password, the SAP HANA administrator must activate and set up the user self-service feature.

Procedure

1. Configure the XSSQLCC technical user required to run the user self-service tools.

A technical user is required to execute user self-service requests; the technical user must be granted a role based on the role template `sap.hana.xs.selfService.user.roles::USSExecutor` and associated with the XSSQLCC artifact `selfService.xssqlcc`.

2. Set the required user-self-service parameters in the `xsengine.ini` file.

As part of the process of enabling user self-service tools in SAP HANA, you must set a number of configuration parameters, for example, to specify the email address to use when responding to user requests or enable support for password-reset services. The parameters must be set in the `user_self_service` section of the `xsengine.ini` file.

i Note

If the section `user_self_service` does not already exist, the SAP HANA administrator must create it.

3. Configure the SMTP server that SAP HANA XS applications can use to send e-mails.

An SMTP server is required to send automatic e-mails in response to the requests users make with SAP HANA user-self-service tools.

i Note

You can configure only one SMTP server per SAP HANA XS server. If an SMTP server is already available, you can skip this step.

4. Configure access to the user self-service administration tools.

You must assign a role based on the role template

`sap.hana.xs.selfService.admin.roles::USSAdministrator` to the user who requires access to the user-self-service administration tools. The user self-service administrator maintains user self-service requests and access blacklists and whitelists.

→ Tip

The user self-service administrator is the user who owns the e-mail address defined in the `sender_email` parameter in the `user_self_service` section of the `xsengine.ini` SAP HANA configuration file.

Related Information

[Set up the Technical User for Self-Service Tools \[page 1597\]](#)

[Configure an SMTP Server for User Self-Service Tools \[page 1598\]](#)

[Configure Access to User-Self-Service Administration Tool \[page 1600\]](#)

12.1.10.1.1 User Self-Service Roles

Dedicated roles are provided to enable access to and the administration of user-self-service tools.

User-self-service tools enable users to request basic database-account services using tools displayed in the user interface. For example, if the self-service tools are enabled, users can request the creation of a new account or a password reset if a password has been forgotten. Additional tools are provided to help administrate the user-self-service requests.

→ Recommendation

As repository roles delivered with SAP HANA can change when a new version of the package is deployed, either do not use them directly but instead as a template for creating your own roles, or have a regular review process in place to verify that they still contain only privileges that are in line with your organization's security policy. Furthermore, if repository package privileges are granted by a role, we recommend that these privileges be restricted to your organization's packages rather than the complete repository. To do this, for each package privilege (`REPO.*`) that occurs in a role template and is granted on `.REPO_PACKAGE_ROOT`, check whether the privilege can and should be granted to a single package or a small number of specific packages rather than the full repository.

User Self-Service Roles

SAP HANA Role	Description
<code>sap.hana.xs.selfService.user.roles::USSAdministrator</code>	<p>Role assigned to the user responsible for administrating the requests sent by users using self-service tools. For example, it provides access to the <i>USS Administration</i> tool, which enables the activation of users who request a new user account in the SAP HANA database and allows the user-self-service administrator to maintain self-service-specific blacklists for user requests, e-mail addresses, domains, and IP addresses.</p> <p>The <i>USS Administrator</i> role also provides access to the tools required to assign roles to (and activate) users in SAP HANA, for example:</p> <ul style="list-style-type: none">• System privileges: USER ADMIN• Object privileges: SELECT on the tables USERS (SYS) and USER_PARAMETERS (SYS)
<code>sap.hana.xs.selfService.user.roles::USSExecutor</code>	<p>Role assigned to the technical user that will be used to respond to and execute user-self-service requests, for example, to create a new account or request a new password.</p>

12.1.10.1.2 Set up the Technical User for Self-Service Tools

Configure the configuration connection (XSSQLCC) and the technical user which are required to execute user self-service requests.

Prerequisites

To complete the steps in this task, you must have the following privileges:

- Access to SAP HANA as the administrator
- Access to specific features provided by the *SAP HANA XS Administration Tool* and the *SAP HANA Web-based Development Workbench*, which requires roles based on the following role templates:
 - `sap.hana.xs.admin.roles::RuntimeConfAdministrator`
 - `sap.hana.xs.admin.roles::SQLCCAdministrator`
 - `sap.hana.xs.ide.roles::SecurityAdmin`

Context

A technical user is required to execute user self-service requests; the technical user must have a role based on the role template `sap.hana.xs.selfService.user.roles::USSExecutor` and associated with the design-time XSSQLCC artifact `selfService.xssqlcc`.

Procedure

1. Create the XSSQLCC technical user required to execute the user self-service requests.
 - a. Open the *SAP HANA Web-based Development Workbench* and start the *Security* tool.
The *Security* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/ide/security`.
- i Note**
Access to the *Security* tool in the *SAP HANA Web-based Development Workbench* requires a role based on the role template `sap.hana.xs.ide.roles::SecurityAdmin`.
- b. Right-click the node **Security > Users** and choose *New User*
 - c. Specify the required details for the new technical user.
You must provide a name and authentication credentials.
 - d. Assign a role based on the role template `sap.hana.xs.selfService.user.roles::USSExecutor` to the new technical user.
 - e. Save your changes to add the new technical user.
2. Assign the new technical user to the `selfService.xssqlcc` artifact.

The technical user you assign to the `selfService.xssqlcc` artifact executes all user-self-service requests, which requires a role based on the role template `sap.hana.xs.selfService.user.roles::USSExecutor`. The `selfService.xssqlcc` artifact provides the appropriate access to SAP HANA.

- a. Start the Web-based *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

To edit `xssqlcc` artifacts with the *SAP HANA XS Administration Tool*, you must have roles based on the following role templates: `sap.hana.xs.admin.roles::RuntimeConfAdministrator` and `sap.hana.xs.admin.roles::SQLCCAdministrator`.

- b. Locate the artifact SQL connection-configuration artifact `selfService.xssqlcc`.

In the *Application Objects* screen, navigate to the package `/sap/hana/xs/selfService/user`.

- c. Assign a technical user to the `selfService.xssqlcc` artifact.

This is the technical user who will be used to execute all user-self-service requests. The user must be assigned a role based on the role template

`sap.hana.xs.selfService.user.roles::USSExecutor`. You must provide the user name and the corresponding password.

Related Information

[Enable User Self-Service Tools \[page 1594\]](#)

12.1.10.1.3 Configure an SMTP Server for User Self-Service Tools

An SMTP server is required to enable SAP HANA to respond to user self-service requests.

Prerequisites

To complete the steps in this task, you must have the following privileges:

- Access to SAP HANA as the administrator
- Access to specific features provided by the *SAP HANA XS Administration Tool* and the *SAP HANA Web-based Development Workbench*, which requires roles based on the following role templates:
 - `sap.hana.xs.admin.roles::RuntimeConfAdministrator`
 - `sap.hana.xs.admin.roles::SMTPDestAdministrator`

Context

To enable SAP HANA to send automatic e-mails in response to the requests users make with SAP HANA user-self-service tools, you must configure a new SMTP server, or make SAP HANA aware of an existing SMTP server.

i Note

You can configure only one SMTP server per SAP HANA XS server. If an SMTP server is already configured, you can use the configured server; you do not have to complete this task.

Procedure

1. Start the Web-based *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

To access to the *SAP HANA XS Administration Tool*, you need a role based on the role template `sap.hana.xs.admin.roles::RuntimeConfAdministrator`.

2. Start the *SMTP Configuration* tool .

i Note

To access to the *SMTP Configuration* tool in the *SAP HANA XS Administration Tool*, you need a role based on the delivered role `sap.hana.xs.admin.roles::SMTPDestAdministrator`.

3. Specify the details of the SMTP server that the user-self-service tools use to reply to service requests. You need to specify the fully qualified domain name of the SMTP server and the port to use for connections, for example, 25 (standard).

→ Tip

For more information about setting up an SMTP server, see *Related Links* below.

Related Information

[Enable User Self-Service Tools \[page 1594\]](#)

[Maintaining SMTP Server Configurations \[page 1568\]](#)

12.1.10.1.4 Configure Access to User-Self-Service Administration Tool

SAP HANA provides an administration tool that enables you to maintain user self-service requests.

Context

Access to the user-self-service administration tools is only possible to users with a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`. The user self-service administrator maintains user self-service requests and access-control blacklists and whitelists.

→ Tip

The user self-service administrator is the user who owns the e-mail address defined in the `sender_email` parameter in the `user_self_service` section of the `xsengine.ini` SAP HANA configuration file.

Procedure

1. Open the *SAP HANA Web-based Development Workbench* and start the *Security* tool.
The *Security* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/ide/security`.

i Note

To access to the *Security* tool in the *SAP HANA Web-based Development Workbench*, you need a role based on the role template `sap.hana.xs.ide.roles::SecurityAdmin`.

2. Configure the user-self-service administrator.
You can create a new user or assign a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator` to an existing user.
 - a. In the *Security* tool, right-click the node **Security > Users** and choose the user for whom you want to enable access to the user-self-service administration tools.
 - b. Assign a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator` to the selected user.
 - c. Save your changes.
3. Log on to the user-self-service administration tool as the new user-self-service administrator.
Verify that you have the permissions required to access to the *USS Administration* tool; the tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/selfService/admin`.

i Note

To log on to the *USS Administration* tool, use the name and password of the user to whom you assigned the role based on the role template

`sap.hana.xs.selfService.admin.roles::USSAdministrator` in the previous step.

Related Information

[Enable User Self-Service Tools \[page 1594\]](#)

[Display all User Self-Service Requests \[page 1604\]](#)

[Maintain User Self-Service Access Lists \[page 1613\]](#)

12.1.10.1.5 Maintain User Self-Service Initialization Parameters

Selected INI parameters can be used to configure how the USS tools respond to user requests and which actions are allowed by default.

Prerequisites

SAP HANA user roles are used to determine the level of access to the features provided by the SAP HANA XS administration tools. To access the tools required to maintain user self-service requests, you must have a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

Context

As part of the process of enabling user self-service tools in SAP HANA, you must set a number of configuration parameters, for example, to activate the self-service tools, specify the email address to use when responding to user requests, or enable support for password-reset services. The parameters you maintain here are synchronized with the corresponding parameters in the `user_self_service` section of the `xsengine.ini` file for the SAP HANA system where you want make self-service tools available.

To display and maintain the initialization parameters for the user self-service tools, perform the following steps:

Procedure

1. Start the user-self-service administration tool.

The *USS Administration* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/selfService/admin`

i Note

To log on, use the name and password of the user who has been assigned a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

2. Display the list of user-self-service requests.

When you open the *USS Administration* tool, the list of user-self-service requests is displayed by default.

i Note

You can also use the *USS Administration* tool to maintain access lists.

3. Choose the *INI Parameters* tool.

4. Set the initialization parameters as required.

Some parameters are enabled (true) or disabled (false); other parameters require a value to be set, for example, a user's e-mail address or the maximum number of times a user can request a new account with USS tools.

i Note

The parameters you maintain here are synchronized with the corresponding parameters in the `user_self_service` section of the `xsengine.ini` file for the SAP HANA system where you want make self-service tools available.

Related Information

[User Self-Service Initialization Parameters \[page 1602\]](#)

12.1.10.1.5.1 User Self-Service Initialization Parameters

Initialization (INI) parameters can be used to configure which USS tools are enabled and how the USS tools react to user requests.

In the *USS Administration* tool, the *INI Parameters* tool displays the mandatory parameters that must be set to enable and configure user-self-service and, in some cases, specify how they can be used. The following table indicates which parameters must be set.

Note

The USS initialization parameters you set with USS administration tools correspond to (and are synchronized with) the SAP HANA parameters listed in the `user_self_service` section of the `xsengine.ini` configuration file.

USS INI Parameter Details

UI Element	Description	Parameter Name	Default
Automatic User Creation	<p>Controls if a user creation request requires approval from user administration. In both cases the administrator has to assign roles to the new user.</p> <ul style="list-style-type: none">Disabled: Requests for a new user account require administrator approval for account activation.Enabled: The user is automatically created and activated as a restricted user.	<code>automatic_user_creation</code>	Disabled/False
Forgot Password	<p>Defines if the system supports password recovery with user-self-service tools. The parameter controls not only the display of the Forgot Password button in the UI logon screen but also the enablement of the corresponding user-self-service backend services.</p>	<code>forgot_password</code>	Disabled/False
Request New user	<p>Enables system support for user-self-service tools. The parameter controls not only the display of the Request New User button in the UI logon screen but also the enablement of the corresponding user-self-service backend services.</p>	<code>request_new_user</code>	Disabled/False
Reset Locked User	<p>Enables support for a password reset for a locked user. Reset password will be forbidden for locked users if the value is Disabled.</p>	<code>reset_locked_user</code>	Disabled/False
Sender E-Mail Address	<p>The email address used for sending out auto-generated replies to user self-service requests, for example, <code>uss.admin@acme.com</code>. Ideally, this is the e-mail address used by the self-service administrator, who is assigned a role based on the role template <code>sap.hana.xs.selfService.admin.roles:USSAdministrator</code> and maintains self-service requests and access lists.</p>	<code>sender_email</code>	None
Token Expiry Time	<p>The time duration (in seconds) for which a generated token (and the corresponding request for a new user or password reset) is valid.</p>	<code>token_expiry_time</code>	3600
User Creation Request Count	<p>The number of times a user can use user-self-service tools to request a new user account. The user is determined by a combination of user name and e-mail address.</p>	<code>user_creation_request_count</code>	3

Optional USS Parameters

It is possible to customize the background image displayed in the logon Web page, for example, by specifying the URL to the image displayed as background in the logon screen. However the following prerequisites apply:

- The image file specified in the URL must be reachable by http(s)
- The URL does not require authentication or authorization
- The recommended minimum resolution of the specified background image is: 1600*1200
- A technical user has to be assigned to the XSSQLCC artifact `/sap/hana/xs/selfService/user/selfService.xssqlcc`. The technical user must also be assigned a role based on the role template `sap.hana.xs.selfService.admin.roles::USSExecutor`. This user will be used to query the details from the server.

Note

The parameter `login_screen_background_image` must be set in the `httpservlet` section of the SAP HANA `xengine.ini` configuration file and can only be set with SAP HANA studio tools.

Optional User Self-Service Configuration Parameters

Parameter Name	Section Name	Description	Example	Default
<code>login_screen_background_image</code>	<code>httpservlet</code>	URL to the image displayed as background in the logon screen	<code>/sap/hana/xs/ui/Image.jpg</code>	None

Related Information

[Maintain User Self-Service Initialization Parameters \[page 1601\]](#)

12.1.10.2 Display all User Self-Service Requests

Display a list of all the user creation requests which have been sent using user self-service tools.

Prerequisites

SAP HANA uses roles to determine the level of access to the features provided by the SAP HANA XS administration tools. To access the tools required to maintain user self-service requests, you must have a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

Context

The user-self-service administrator is the user associated with the e-mail address defined in the `xseengine.ini` parameter `sender_email`. The user-self-service administrator can use the *USS Administration* tool to view a list of all the self-service requests received from users. Each user self-service request includes the following details:

- User name
- Creation date and time
- Number of pending self-service requests made by the same user

To display all user self-service requests, perform the following steps:

Procedure

1. Start the user-self-service administration tool.

The *USS Administration* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/selfService/admin`.

i Note

To log on, use the name and password of the user who has been assigned a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

When you open the *USS Administration* tool, the list of user-self-service requests is displayed by default.

2. Display the access-control list that you want to maintain.

You can maintain access-control lists for the following conditions:

- Domain
- E-mail address
- IP ranges

3. Maintain entries for the selected access-control list.

You can use the *Add* and *Delete* buttons to manage the list entries.

→ Tip

To delete an entry from an access-control list, first check one or more items in the list and choose *Delete*.

Related Information

[Display all User Self-Service Requests \[page 1604\]](#)

[Activate a User Account \[page 1610\]](#)

[Reject a User Self-Service Request \[page 1611\]](#)

12.1.10.3 Request a New User Account

Request a new user account with user-self-service tools.

Prerequisites

- User-self-service tools are enabled in SAP HANA
- The required technical user (with the role *USSExecutor* is configured and available to respond to user-self-service requests

Context

If the self-service tools are enabled, a user can use the tools to request a new user account in the SAP HANA database. A valid e-mail address is required to complete the account-creation process, and the administrator must activate the new account and assign user roles and privileges.

To request a new database account in SAP HANA, a user must perform the following steps:

Procedure

1. Logon to SAP HANA using the Web-based interface.
The SAP HANA *Logon* screen is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/formLogin/login.html`.
2. Request a new user SAP HANA account.
Click the *Request Account* link in the bottom right-hand side of the logon screen.
3. Specify basic details for the new user.
In the *Request Account* screen, you must supply a name and a valid e-mail address, which the user-self-service tools use to respond to the request.
 - a. Enter a name for the new database user.
 - b. Enter a valid e-mail address for the new database user.
The e-mail address is used to sent the user messages with links to use to start the account activation process.
4. Submit the request for a new account.
After submitting the account-creation request, the user receives the following automatically generated e-mails:
 - Address verification
An e-mail with a link that verifies the target e-mail address
 - User-self-service request administratration
An e-mail that contains the following links:

- Open the *SAP HANA XS Administration Tool* tool that enables an account be set up and activated for the new user
 - Display a list of all pending user-self-service requests
5. Set a password and security question for the new user account
The user requesting the new database account must set a password and choose a security question that is used in the event of a forgotten-password request. An answer must be supplied for the selected security question.
 6. Activate the new user account.
The user self-service administrator must activate the new user account to enable the new user to log on to SAP HANA. Activation involves assigning roles to the new user as well as privileges, for example: objects, application, package.

12.1.10.4 Maintain Your User Profile

Each user account is associated with a profile; the user who owns the profile must adjust the settings to suit personal preferences.

Prerequisites

- User-self-service tools are enabled in SAP HANA.
- A user profile exists; a user profile is created automatically on activation of a user account in SAP HANA.
- The profile owner has the privileges granted by the role `sap.hana.xs.formLogin.profile::ProfileOwner`.

Context

When a new user account is activated, the corresponding account profile is created with default settings. The new user must log on to SAP HANA and adjust some of the default settings, for example, the default password. It is also mandatory to choose a security question and set the corresponding answer.

Procedure

1. Log on to SAP HANA using the Web-based interface.
The SAP HANA *Logon* screen is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/formLogin/login.html`.
2. Start the profile manager.
The *Manage Profile* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/formLogin/profile/`.

3. Maintain your security settings.

It is mandatory to choose a security question from the drop-down list and provide a corresponding answer. You can also change the e-mail address to use for communication.

i Note

The question and answer are used to confirm logon credentials, whenever the owner of the profile attempts to make any changes to the profile.

4. Change the initial (default) password.

i Note

In SAP HANA, rules apply that restrict which characters you can use in the password you set.

5. Maintain your profile preferences.

You can change set the language locale for your account and set preferences for the way that the date and time is displayed, for example:

- *Date Format: YYYY-MM-DD* ("2014-12-25")
- *Time Format:*
 - *HH24:MI* ("15:30")
 - *HH12:MI* ("3:30pm")
- *Locale: English (en)*

i Note

Application developers need to ensure that the applications they create are able to take account of the preference set in a user's profile.

Related Information

[Request a New User Account \[page 1606\]](#)

12.1.10.4.1 User Profile Details

Each user account has a corresponding account profile.

When a new user account is activated, the corresponding account profile is created with default settings. The new user must log on to SAP HANA and adjust some of the default settings, for example, the default password. It is also mandatory to choose a security question and set the corresponding answer. The *User Self Services Manage Profile* tool displays the following screens to help you maintain details of the SAML service provider:

- [Security Settings \[page 1609\]](#)
- [Preferences \[page 1609\]](#)
- [Change Password \[page 1609\]](#)

Security Settings

The *Security Settings* screen area in the USS *Manage Profile* tool enables you to maintain details of the security settings for your SAP HANA user account. The following table indicates which details can be maintained.

UI Element	Description	Example
<i>Email Address</i>	The e-mail address of the user to whom the account and profile belong. USS notifications are sent to the specified address.	Kwame.Ampomah@acme.com
<i>Security Question</i>	The security question to ask when you make any changes to the user profile details.	What is your favorite sport?
<i>Security Answer</i>	Text string that you use as the answer to the security question	squash

Preferences

The *Preferences* screen area in the USS *Manage Profile* tool enables you to maintain details of the display preferences for your SAP HANA user account. The following table indicates which details can be maintained.

UI Element	Description	Example
<i>Date Format</i>	The way in which the date is displayed in the applications you use, for example, <i>2014-12-25</i>	YYYY-MM-DD
<i>Time Format</i>	The way in which the time is displayed in the applications you use, for example, <i>15:30</i> (HH24:MI) or <i>3:30pm</i> (HH12:MI)	HH24:MI
<i>Locale</i>	The language environment and settings to apply for the applications you use	English (en) or Chinese (zh)

Change Password

The *Change Password* screen area in the USS *Manage Profile* tool enables you to maintain details of your SAP HANA user account. The following table indicates which details can be maintained.

UI Element	Description	Example
<i>Old Password</i>	The initial password assigned when the account was activated or, if changed, the currently valid password	*****
<i>New Password</i>	The new password	*****
<i>Repeat Password</i>	Confirm the new password you entered in <i>New Password</i>	*****

Related Information

[Maintain Your User Profile \[page 1607\]](#)

12.1.10.5 Activate a User Account

Enable a new user account in the SAP HANA database in response to a user self-service request.

Prerequisites

SAP HANA uses roles to determine the level of access to the features provided by the SAP HANA XS administration tools. To access the tools required to maintain user self-service requests, you must have a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

Context

When a user requests a new account, the user account is created but disabled by default. The user who sent the request cannot use the account to log on to SAP HANA until the SAP HANA user-self service administrator activates the account. The self-service administrator must manually activate the account and assign the necessary roles, too.

i Note

On activation of the new user account, an e-mail is automatically sent to the user containing the security token required to enable the new user to set a password for the new account.

To activate a new account in response to a user self-service request, perform the following steps:

Procedure

1. Start the user-self-service administration tool.

The *USS Administration* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/selfService/admin`.

i Note

To log on, use the name and password of the user who has been assigned a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

2. Display the list of user-self-service requests.

When you open the *USS Administration* tool, the list of user-self-service requests is displayed by default.

i Note

You can also use the *USS Administration* tool to maintain access lists.

3. Assign roles to a new user.
 - a. In the *Username* column of the *User Self Service Requests* screen, check the box next the user you want to activate.
 - b. In the *Administration* column, choose *Assign Roles*.

The link opens the *Security* tool in the SAP HANA Web-based Development Workbench and displays the selected user. Select the appropriate roles to assign to the new user from the list of roles displayed.

i Note

To help decide which roles are appropriate for the user request, use the path indicated in the *Request Origin* column to see which tool the user is trying to access. For example, `/sap/hana/ide/editor` is the SAP HANA Editor tool, which requires a role based on the role template `sap.hana.ide.roles::EditorDeveloper`.

4. Activate the selected new user.

In the *User Self Service Requests* page, choose *Activate and Notify* to send an e-mail to the corresponding user indicating that the requested account is active and ready for use.

Related Information

[Reject a User Self-Service Request \[page 1611\]](#)

[Maintain User Self-Service Access Lists \[page 1613\]](#)

[Display all User Self-Service Requests \[page 1604\]](#)

12.1.10.6 Reject a User Self-Service Request

Refuse a self-service request to create a new user account in the SAP HANA database.

Prerequisites

SAP HANA uses roles to determine the level of access to the features provided by the SAP HANA XS administration tools. To access the tools required to maintain user self-service requests, you must have a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

Context

When a user requests a new account, the user account is created but disabled by default. The user who sent the request cannot use the account to log on to SAP HANA until the SAP HANA user-self service administrator activates the account and assigns the appropriate roles. The user-self-service administrator can also choose to reject the request for a new user account in the SAP HANA database, for example, by adding the user to the user-requests blacklist.

i Note

On activation of the new user account, an e-mail is automatically sent to the user containing the security token required to enable the new user to set a password for the new account.

To refuse a self-service request to create a new user account in the SAP HANA database, perform the following steps:

Procedure

1. Start the user-self-service administration tool.

The *USS Administration* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/selfService/admin`.

i Note

To log on, use the name and password of the user who has been assigned a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

2. Display the list of user-self-service requests.

When you open the *USS Administration* tool, the list of user-self-service requests is displayed by default.

i Note

You can also use the *USS Administration* tool to maintain access lists.

3. Reject the user's request for a new database account.
 - a. In the *Username* column of the *User Self Service Requests* screen, check the box next the user, whose request for a new account you want to reject.
 - b. Choose *Add to blacklist* in the bottom right-hand corner of the screen.

The link opens the *Security* tool in the SAP HANA Web-based Development Workbench and displays the selected user. Select the appropriate roles to assign to the new user from the list of roles displayed.
4. Check the rejected user has been added to the user-requests blacklist.

Related Information

[Maintain User Self-Service Access Lists \[page 1613\]](#)

12.1.10.7 Maintain User Self-Service Access Lists

Access to self-service tools can be controlled using blacklists and whitelists, for example, for email addresses.

Prerequisites

SAP HANA uses roles to determine the level of access to the features provided by the SAP HANA XS administration tools. To access the tools required to maintain user self-service requests, you must have a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

Context

The user self-service administrator can control access to self-service features by maintaining blacklists and whitelists for the following areas:

- User requests
- Network domains
- IP addresses
- E-mail addresses
- DB users

i Note

Users whose requests exceed the value set in the `xsengine.ini` parameter `user_creation_request_count` are no longer able to submit any requests. If necessary, the administrator can add such users to the access blacklist.

To display all user self-service access lists, perform the following steps:

Procedure

1. Start the user-self-service administration tool.

The *USS Administration* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/selfService/admin`.

i Note

To log on, use the name and password of the user who has been assigned a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

2. Display the list of user-self-service requests.

When you open the *USS Administration* tool, the list of user-self-service requests is displayed by default.

i Note

You can also use the *USS Administration* tool to maintain access lists.

Related Information

[User Self-Service Access Lists \[page 1614\]](#)

12.1.10.7.1 User Self-Service Access Lists

Access to self-service features is controlled by blacklists and whitelists.

The user self-service administrator can control access to user-self-service tools using the *USS Administration* by maintaining access lists. The access lists included with the *USS Administration* are described in the following table.

User Self-Service Access List Details

List Name	Description
User Requests	A list of all pending requests sent with user-self-service tools, including the name of the user who sent the request and the corresponding e-mail address. Users who have more requests than the value set in the <code>xengine.ini</code> parameter <code>user_creation_request_count</code> are automatically added to the access blacklist.
Network Domains	A list of network domains, which can be used to permit or deny user self-service requests from one or more specific domains, for example, "acme.com". If a user self-service request for a new user account arrives from a user with an e-mail address associated with a whitelisted domain, the new user account is created as a restricted user and activated without requiring any administrator intervention. Users on the domain black list are no longer permitted to create a user self-service request.
IP Addresses	A list of IP addresses (or names), which can be used to permit or deny user self-service requests from one or more specific IP addresses, for example, "* .122 .10". The same rules for blacklists and whitelists apply as for network domains above.
E-mail addresses	A list of e-mail addresses, which can be used to permit or deny user self-service requests from a specific e-mail address, for example, "joe.doe@acme.com" or "jane.doe@acme.com". The same rules for blacklists and whitelists apply as for network domains and IP addresses above.

List Name	Description
DB Users	<p>The names of the database users who are not allowed to change their respective SAP HANA password using USS reset-password tools, for example, j_oedoe or j_anedoe. The following additional restrictions apply:</p> <ul style="list-style-type: none"> • By default, it is not possible to use USS tools to reset the password for the SYSTEM user. • The USS administrator cannot add to the <i>DB Users</i> list any user who logs on to SAP HANA with single sign-on (SSO) credentials. • Users who log on to SAP HANA with SSO credentials cannot use USS tools to reset their password.

12.1.10.8 Maintain User Self-Service E-Mail Templates

Default templates enable you to format the contents of the auto-generated e-mails sent when user self-service (USS) tools are employed to request a new account in SAP HANA or recovery a forgotten password.

Prerequisites

SAP HANA uses roles to determine the level of access to the features provided by the SAP HANA XS administration tools. To access the tools required to maintain user self-service requests, you must have a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

Context

The user self-service administrator can modify the contents of the automatically generated e-mails that are sent to users during the USS account-creation process. Templates exist for the responses to the following actions: user requests, account activation, and forgotten passwords.

To display and maintain the current e-mail templates for user self-service features, perform the following steps:

Procedure

1. Start the user-self-service administration tool.

The *USS Administration* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/selfService/admin`

i Note

To log on, use the name and password of the user who has been assigned a role based on the role template `sap.hana.xs.selfService.admin.roles::USSAdministrator`.

2. Display the list of user-self-service requests.

When you open the *USS Administration* tool, the list of user-self-service requests is displayed by default.

i Note

You can also use the *USS Administration* tool to maintain access lists.

3. Choose the *Email Templates* tool.
4. Choose the required e-mail template.

The following templates are available for automatically generated e-mails:

- User request
E-mail sent in response to a user self-service request for a new SAP HANA user account
- User activation
E-mail sent when a new SAP HANA user account has been activated
- Password Recovery
E-mail sent in response to a user self-service request to set a new SAP HANA password, for example, because the user has forgotten the current password.

Related Information

[User Self-Service E-Mail Templates \[page 1616\]](#)

12.1.10.8.1 User Self-Service E-Mail Templates

USS provides templates that can be used to format the content of auto-generated e-mails.

In the *USS Administration* tool, the *Email Templates* tool displays information about the templates used to format the content of auto-generated e-mails that are used during the process of creating a new SAP HANA user account. You can use the *Email Templates* tab to maintain the following details:

- [User Request \[page 1616\]](#)
- [User Activation \[page 1617\]](#)
- [Forgot Password \[page 1617\]](#)

User Request

The *User Request* tab in the *Email Templates* tool enables you to maintain templates that are used to generate the e-mails sent in response to a user request to create a new account in SAP HANA; the e-mails are sent to the USS administrator and the user who submitted a USS request. The following table indicates which information can be viewed and modified.

User Request E-Mail Template Details

UI Element	Description	Example
<i>To</i>	The email address of the USS administrator	admin.uss@acme.com
<i>Subject</i>	The text you want to appear in the e-mail's <i>Subject</i> box	New user account
<i>Body</i>	The text of the e-mail sent either to the USS admin indicating that a new request for a SAP HANA user account has been received and needs attention or to the user who submitted a request and indicating that the request for a new account has been received and is being processed	Dear USS Admin, ...

User Activation

The *User Activation* tab in the *Email Templates* tool enables you to maintain templates for the account-activation e-mails sent to the user who uses USS tools to submit a request for a new account in SAP HANA; the e-mail informs the user that the requested account is active and can be used to log on to SAP HANA. The following table indicates which information can be viewed and modified.

User Activation E-Mail Template Details

UI Element	Description	Example
<i>To</i>	The email address of the user whose new accounts has been activated	admin.uss@acme.com
<i>Subject</i>	The text you want to appear in the e-mail's <i>Subject</i> box	SAP HANA account status
<i>Body</i>	The text of the e-mail sent either to the new SAP HANA user indicating that an account has been activated and can be used to log on to SAP HANA	Dear [<i><User Name></i>], ...

Forgot Password

The *Forgot Password* tab in the *Email Templates* tool enables you to maintain the template used to generate e-mails that are sent to SAP HANA users who submit a USS request to reset a password. The following table indicates which information can be viewed and modified.

Forgot Password E-Mail Template Details

UI Element	Description	Example
<i>To</i>	The email address of the user who submitted a request to reset a password	jane.doe@acme.com
<i>Subject</i>	The text you want to appear in the e-mail's <i>Subject</i> box	Reset account password
<i>Body</i>	The text of the e-mail to the SAP HANA user indicating that a request to reset an SAP HANA password has been received and action is required from the user	Dear [<i><User Name></i>], ...

Related Information

[Maintain User Self-Service E-Mail Templates \[page 1615\]](#)

12.1.11 Scheduling XS Jobs

Scheduled jobs define recurring tasks that run in the background. The JavaScript API `$.jobs` allows developers to add and remove schedules from such jobs.

If you want to define a recurring task, one that runs at a scheduled interval, you can specify details of the job in a `.xsjob` file. The time schedule is configured using `cron`-like syntax. You can use the job defined in an `.xsjob` file to run an XS Javascript or SQLScript at regular intervals. To create and enable a recurring task using the `xsjob` feature, you perform the following high-level tasks:

i Note

The tasks required to set up a scheduled job in SAP HANA XS are performed by two distinct user roles: the application developer and the SAP HANA administrator. In addition, to maintain details of an XS job in the [SAP HANA XS Administration Tool](#), the administrator user requires the privileges granted by the role template `sap.hana.xs.admin.roles::JobAdministrator`.

Setting up Scheduled Jobs in SAP HANA XS.

Step	Task	User Role	Tool
1	Create the function or script you want to run at regular intervals	Application developer	Text editor
2	Create the job file <code>.xsjob</code> that defines details of the recurring task	Application developer	Text editor
3	Maintain the corresponding runtime configuration for the <code>xsjob</code>	SAP HANA administrator	XS Job Dashboard
4	Enable the job-scheduling feature in SAP HANA XS	SAP HANA administrator	XS Job Dashboard
5	Check the job logs to ensure the job is running according to schedule.	SAP HANA administrator	XS Job Dashboard

Related Information

[The XSJob File \[page 1630\]](#)

[Tutorial: Schedule an XS Job \[page 1627\]](#)

12.1.11.1 Maintain XS Job Details

XS job schedules are defined by developers; the XS job-scheduling feature must be set up by a system administrator.

Prerequisites

To enable the XS Job schedule feature in SAP HANA XS, the following prerequisites apply:

- You have administrator access to an SAP HANA system.
- You have been granted a role based on the role template `sap.hana.xs.admin.roles::JobAdministrator`.
- An XS job file that has been activated in the repository.

Context

To enable developers to define and deploy job schedules using the XS job feature, the system administrator must first set up the environment and enable some essential options.

Procedure

1. Enable the job-scheduling feature in SAP HANA XS.

This step requires the permissions granted to the SAP HANA administrator.

i Note

It is not possible to enable the scheduler for more than one host in a distributed SAP HANA XS landscape.

- a. In the *XS Job Dashboard* set the *Scheduler Enabled* toggle button to **YES**.

toggling the setting for the *Scheduler Enabled* button in the *XS Job Dashboard* also changes the current value of the SAP HANA configuration variable `xsengine.ini > scheduler > enabled`, which is set in the *Configuration* tab of the SAP HANA studio's *Administration* perspective.

2. Maintain the XS job's runtime configuration.

- a. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

- b. Open the *XS Job Dashboard*.

i Note

To maintain details of an XS job using the Web-based *XS Administration Tool* you need the privileges granted in the SAP HANA user role `sap.hana.xs.admin.roles::JobAdministrator`.

- c. Maintain the details of the XS job.

In the *Job Details* tab, select the XS Job whose details you want to maintain. In the *Configuration* tab, you need to specify the following details:

- o *User*
The user account in which the `xscron` job runs, for example, **SYSTEM**
- o *Password*
For security reasons, you must provide a password for the specified user.

i Note

If you do not provide a user password, you cannot save the changes to the XS Job object's runtime configuration.

- o *Locale*
The language encoding required for the locale in which the `xscron` job runs, for example, **en_US**
- o *Start/Stop time*
An optional value to set time during which the `xscron` job runs. You must enter the values using the syntax used for the SAP HANA data type `LocalDate` and `LocalTime`, for example, **2013-11-05 00:30:00** (thirty minutes past midnight on the 5th of November 2013).
- o *Active*
Enable or disable the job schedule

- d. Save the job.

Choose *Save Job* to save and activate the changes to the job schedule.

3. Check the logs to ensure the job is running according to schedule.

You can view the list of `xsjob` schedules in the *Job Details* tab of the *XS Job Details* window. The information displayed includes the XS cron setup that defines the schedule, the current status of the job schedule, as well as the start and finish times.

Related Information

[The XS Job Dashboard \[page 1620\]](#)

[The XS Job File \[page 1630\]](#)

12.1.11.1.1 The XS Job Dashboard

The *XS Job Dashboard* is the central point of control for monitoring and maintaining job schedules that have been defined using the XS Job syntax.

The *XS Job Dashboard* displays details of the currently active job schedules that have been configured for the selected SAP HANA system using XS job files. The XS job file uses a cron-like syntax to specify the schedule at

which the service defined in an XS JavaScript or SQLScript must run. You can use the *Scheduler Enabled* button in the *XS Job Dashboard* to enable schedules for all XS jobs globally.

Note

Toggling the setting for the *Scheduler Enabled* button also changes the current value of the SAP HANA configuration variable `xsengine.ini > scheduler > enabled`, which is set in the *Configuration* tab of the SAP HANA studio's *Administration* perspective.

For each XS job displayed in the *XS Job Dashboard*, you can see the following details:

- **Name**
The name of the XS Job; this is name of the design-time artifact in the SAP HANA repository, for example, `MyJob.xsjob`
- **Package**
The name of the repository package that contains the XS Job
- **User**
The name of the user whose database account is used to run the XS Job schedule
- **Status**
The current status of the XS job schedule, for example, *ACTIVE/INACTIVE*; you can change the status in the *XS Job Details* screen
- **Start/Stop time**
An optional value to set the period of time during which the job runs. You must enter the values using the syntax used for the SAP HANA data type `LocalDate` and `LocalTime`, for example, `2014-11-05 00:30:00` (thirty minutes past midnight on the 5th of November 2014).
- **Session Timeout(s)**
The number of times that the scheduled job run encountered a session timeout
- **Last Run Status**
The status of the scheduled job when it last ran, for example: *Success*, *Error*, or *Running*

Related Information

[Maintain XS Job Details \[page 1619\]](#)

12.1.11.1.2 XS Job Details

Details of the runtime configuration of XS Job schedules and the XS jobs the schedules are used to manage.

In the *XS Job Dashboard*, the *XS Job Details* tab displays information about the currently active job schedules that have been configured for the selected SAP HANA system and the corresponding XS job files. You can use the *XS Job Details* tab to maintain the following details of the XS Jobs' runtime configuration:

- [General Job Details \[page 1622\]](#)
- [Runtime Configuration \[page 1623\]](#)
- [Log Cleanup \[page 1623\]](#)

Job Details

The *Job Details* tab in the *XS Job Details* tool enables you to view details of the XS Jobs that you have defined and scheduled to run, for example, the name of the XS job and a short description. The following table indicates which information can be viewed.

Note

The details displayed are defined in the design-time artifact that describes the selected XS Job.

Job Details

UI Element	Description	Example
<i>Name</i>	Text string used to specify the name (including full repository path) of the XS Job scheduled to run.	sap.hana.testtools::schedule
<i>Description</i>	A short description of the XS job defined in <i>Name</i>	Run XSUnit
<i>Action</i>	Text string used to specify the path to the function to be called as part of the XS Job defined in <i>Name</i>	sap.hana.testtools:TestRunner.xsjs::run

Runtime schedules for XS Jobs contain the following details.

Note

Some of the values described (for example, *Origin* or *Changed ...*) are read only; it is not possible to modify them.

Schedules

UI Element	Description	Example
<i>ID</i>	The ID allocated to the job schedule	3
<i>XCron</i>	The schedule for the specified task (defined in the "action" keyword); the schedule is defined using cron-like syntax.	2015 * * fri 12 0 0
<i>Parameter</i>	A value to be used during the action operation. You can add as many parameters as you like as long as they are mapped to a parameter in the function itself.	Depends on job
<i>Planned Time</i>	The time at which an XS job is expected to run; if it does not run as planned, it is added to the job queue.	2014-11-05 00:30:00
<i>Status</i>	Indicates if the schedule is active or inactive	Active
<i>Start Time</i>	An optional value signifying the beginning of the period of time (schedule) during which the XS job runs	2014-11-05 00:30:00
<i>Finish Time</i>	An optional value signifying the end of the period of time (schedule) during which the XS job runs	2014-11-12 00:30:00
<i>Time Taken (s)</i>	The amount of time taken (in seconds) for the job/action to complete	5

UI Element	Description	Example
<i>Description</i>	A short (optional) description of the XS job schedule.	gfn test schedule
	<p>→ Tip</p> <p>It is not possible to show the <i>Description</i> of a deleted job schedule, even if a description was defined when configuring the original job schedule. If the <i>Description</i> field in the job <i>Logs</i> view is empty, either no description was provided for the corresponding job schedule or the job schedule has been deleted. If no description was provided when configuring a job schedule, the <i>Description</i> field in the <i>Job Details</i> tab always remains empty.</p>	
<i>Origin</i>	The type of object used to define the schedule: <i>DESIGNTIME</i> (repository artifact) or <i>RUNTIME</i> (catalog object).	DESIGNTIME
<i>Changed By</i>	Name of the SAP HANA user who added or changed the XS job schedule	johndoe
<i>Changed At</i>	Time at which the schedule was changed	2015-01-30 14:19:59

Configuration

The *Configuration* tab in the *XS Job Details* tool enables you to maintain details of the runtime configuration for XS Jobs that you have scheduled to run. The following table indicates which information can be maintained.

XS Job Configuration

UI Element	Description	Example
<i>User</i>	The user account in which the xs cron job runs.	SYSTEM
<i>Password</i>	Password for the specified <i>user</i>	****
<i>Locale</i>	The language encoding required for the locale in which the xs cron job runs	en_US
<i>Start Time</i>	Start time for the XS Job using the syntax required by the SAP HANA data type <code>LocalDate</code> and <code>LocalTime</code>	2013-11-05 00:30:00
<i>End Time</i>	End time for the XS Job using the syntax required by the SAP HANA data type <code>LocalDate</code> and <code>LocalTime</code>	2013-11-05 00:30:00
<i>Session Timeout</i>	Time in seconds for which the session is valid	0
<i>Active</i>	Indicates if the schedule is active or inactive	Active

Log Cleanup

The *Log Cleanup* tab in the *XS Job Details* tool enables you to create an XS Job that cleans up the logs of all XS Job currently running in the system. You can also create one schedule for each job in the system and allow users to configure the schedule in the *Job Details* dialog.

By default, XS Job logs are not cleaned up; no logs or log entries are deleted. If a cleanup of XS Job logs is required, the parameters can be set so that only those job-log entries for an XSJob that are older than N days are deleted, where N can be configured as a job parameter. Users can also specify the frequency of the cleanup schedule. The following table indicates which information can be maintained.

! Restriction

To enable or disable the cleanup of XS Job logs, you need the permissions assigned to the *JobAdministrator* role.

XS Job Log Cleanup

UI Element	Description	Example
<i>Enabled</i>	Enable the log-cleanup schedule	Yes
<i>XCron</i>	The schedule for the specified XS Job log-cleanup task; the schedule is defined using cron-like syntax. In this example, the cleanup is scheduled to run every last Sunday of the month at 09:00 hours. (9am)	* * * -1.sun 9 0 0
<i>Day</i>	The number of days for which logs are retained (not cleaned up). For example, 1 retains all XS job logs from the day before the schedule starts and deletes all job logs that are two days old or older.	1

Related Information

[SAP HANA XS Classic Administration Roles \[page 1529\]](#)

[Scheduling XS Jobs \[page 1618\]](#)

12.1.11.2 Clean up XS Job Logs

Clean up the log entries generated in the SAP HANA database by the XS jobs that are running in the SAP HANA system.

Prerequisites

To enable the XS Job schedule feature in SAP HANA XS, the following prerequisites apply:

- You have administrator access to an SAP HANA system.
- You have been granted a role based on the role template `sap.hana.xs.admin.roles::JobAdministrator`.
- You have enabled the job-scheduling feature in SAP HANA XS.

- You have maintained details of the XS Job whose log entries you want to clean up.
- You have enabled the XS Job `sap.hana.xs.admin.jobs.server.common::cleanJobLog` that is used to clean up job-log entries
- You have activated the SQLCC artifact `sap.hana.xs.admin.jobs.server.common::cleanJobLog.xssqlcc` that is used by the cleanup job; this artifact creates a connection to SAP HANA with the `JobLogAdmin` privileges required to remove entries from the XS-job log (as defined in `cleanJobLog`)

Context

XS jobs write their logs to the table `_sys_xs.job_log` in the SAP HANA database. Since this table can grow in size very quickly, as more and more jobs and schedules are created, it is recommended to clean up the old job log entries. You can set up an XS Job that runs at a defined schedule and deletes all old log file entries for a particular XS job from the SAP HANA XS job-log table.

Procedure

1. Maintain the XS job's runtime configuration.

- a. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

- b. Open the *XS Job Dashboard*.

i Note

To maintain details of an XS job using the Web-based *XS Administration Tool* you need the privileges granted in the SAP HANA user role `sap.hana.xs.admin.roles::JobAdministrator`.

2. Configure details of the XS job schedule.

In the *Job Details* tab, select the XS Job whose details you want to maintain. In the *Configuration* tab, you need to specify the following details:

- *User*

The user account in which the `xscron` job runs, for example, **SYSTEM**

- *Password*

For security reasons, you must provide a password for the specified user.

i Note

If you do not provide a user password, you cannot save the changes to the XS Job object's run-time configuration.

- *Locale*

The language encoding required for the locale in which the `xscron` job runs, for example, **en_US**

- *Start/Stop time*

An optional value to set time during which the `xscron` job runs. You must enter the values using the syntax used for the SAP HANA data type `LocalDate` and `LocalTime`, for example, **2013-11-05 00:30:00** (thirty minutes past midnight on the 5th of November 2013).

- *Active*
Enable or disable the job schedule.
3. Ensure that the old log entries written by the XS job are cleaned up.

To enable a scheduled clean up of log entries in the SAP HANA database, you need to set up the following details:

- *Enabled*
Set the status of the job schedule used to clean up the XS job-related log entries
 - *XSCron*
Define the schedule using XS cron syntax (year, month, day, day of the week, hour, minute, second) at which the cleanup job runs.
*** * * -1.sun 9 0 0**
This example runs the job on the last Sunday of every month at 9am.
 - *Day*
Specify the number of days for which log entries should be **retained**. For example, to delete all log entries that are older than two days, enter the value "2".
4. Save the job.
Choose *Save Job* to save and activate the changes to the job schedule.
 5. Check the status of the new job and schedule.
You can view the list of `xsjob` schedules in the *Job Details* tab of the *XS Job Details* window. The information displayed includes the XS cron setup that defines the schedule, the current status of the job schedule, as well as the start and finish times.
 6. Check the logs to ensure the job is running according to schedule.

→ Tip

If the *Description* field for a specific job log is empty in the *View Logs* list, this is typically an indication that either no description was defined for the job or the log has been deleted as part of the cleanup operation.

Related Information

[Maintain XS Job Details \[page 1619\]](#)

[The XS Job Dashboard \[page 1620\]](#)

[XS Job Details \[page 1621\]](#)

12.1.11.3 Tutorial: Schedule an XS Job

The `xsjob` file enables you to run a service (for example, an XS JavaScript or an SQLScript) at a scheduled interval.

Prerequisites

- You have access to an SAP HANA system.
- You have a role based on the role template `sap.hana.xs.admin.roles::JobAdministrator`.
- You have a role based on the role template `sap.hana.xs.admin.roles::HTTPDestAdministrator`.

Note

This tutorial combines tasks that are typically performed by two different roles: the application developer and the database administrator. The developer would not normally require the privileges granted to the `sap.hana.xs.admin.roles::JobAdministrator` role, the `sap.hana.xs.admin.roles::HTTPDestAdministrator` role, or the SAP HANA administrator.

Context

In this tutorial, you learn how to schedule a job that triggers an XS JavaScript application that reads the latest value of a share price from a public financial service available on the Internet. You also see how to check that the XS job is working and running on schedule.

To schedule an XS job to trigger an XS JavaScript to run at a specified interval, perform the following steps:

Procedure

1. Create the application package structure that contains the artifacts you create and maintain in this tutorial. Create a root package called `yahoo`. You use the new `yahoo` package to contain the files and artifacts required to complete this tutorial.

```
/yahoo/  
  .xsapp           // application descriptor  
  yahoo.xsjob      // job schedule definition  
  yahoo.xshttpdest // HTTP destination details  
  yahoo.xsjs       // Script to run on schedule
```

2. Write the XS JavaScript code that you want to run at the interval defined in an XS job schedule. The following XS JavaScript connects to a public financial service on the Internet to check and download the latest prices for stocks and shares.

Create an XS JavaScript file called `yahoo.xsjs` and add the code shown in the following example:

```
function readStock(input) {
```

```

var stock = input.stock;

var dest = $.net.http.readDestination("yahoo", "yahoo");
var client = new $.net.http.Client();
var req = new $.web.WebRequest($.net.http.GET, "/d/quotes.csv?f=a&s=" +
stock);
client.request(req, dest);
var response = client.getResponse();
var stockValue;
if(response.body)
    stockValue = parseInt(response.body.asString(), 10);
var sql = "INSERT INTO stock_values VALUES (NOW(), ?)";
var conn = $.db.getConnection();
var pstmt = conn.prepareStatement(sql);
pstmt.setDouble(1, stockValue);
pstmt.execute();
conn.commit();
conn.close();
}

```

Save and activate the changes in the SAP HANA Repository.

i Note

Saving a file in a shared project automatically commits the saved version of the file to the repository. To explicitly commit a file to the repository, right-click the file (or the project containing the file) and choose **Team > Commit** from the context-sensitive popup menu.

3. Create an HTTP destination file using the wizard to provide access to the external service (via an outbound connection).

Since the financial service used in this tutorial is hosted on an external server, you must create an HTTP destination file, which provides details of the server, for example, the server name and the port to use for HTTP access.

i Note

To maintain the runtime configuration details using the Web-based *XS Administration Tool* you need the privileges granted in the SAP HANA user role `sap.hana.xs.admin.roles::HTTPDestAdministrator`.

Create a file called `yahoo.xshttpdest` and add the following content:

```

host = "download.finance.yahoo.com";
port = 80;

```

Save and activate the changes in the SAP HANA Repository.

4. Create the XS job file using the wizard to define the details of the schedule at which the job runs.

The XS job file uses a `cron`-like syntax to define the schedule at which the XS JavaScript must run. This job file triggers the script `yahoo.xsjs` on the 59th second of every minute and provides the name "SAP.DE" as the parameter for the stock value to check.

Create a file called `yahoo.xsjob` and add the following code:

```

{
  "description": "Read stock value",
  "action": "yahoo:yahoo.xsjs::readStock",
  "schedules": [
    {
      "description": "Read current stock value",
      "xcron": "* * * * * 59",
    }
  ]
}

```

```

        "parameter": {
            "stock": "SAP.DE"
        }
    ]
}

```

Save and activate the changes in the SAP HANA Repository.

5. Maintain the XS job's runtime configuration.

You maintain details of an XS Job's runtime configuration in the *XS Job Dashboard*.

a. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

b. Maintain the details of the XS job.

Note

To maintain details of an XS job using the Web-based *XS Administration Tool* you need the privileges granted in the SAP HANA user role `sap.hana.xs.admin.roles::JobAdministrator`.

You need to specify the following details:

- *User*
The user account in which the job runs, for example, **SYSTEM**
- *Password*
The password required for user, whose account is used to run the job.
- *Locale*
The language encoding required for the locale in which the job runs, for example, **en_US**
- *Start/Stop time*
An optional value to set the period of time during which the job runs. Enter the values using the syntax used for the SAP HANA data type `LocalDate` and `LocalTime`, for example, **2014-11-05 00:30:00** (thirty minutes past midnight on the 5th of November 2014).
- *Active*
Enable or disable the job schedule
- *Session timeout*
Specify the session timeout for this XSJob in seconds. If you specify a value of 0 (zero) seconds for the XSJob's session timeout, the XSJob checks if a value is defined for the `sessiontimeout` key in the section `scheduler` of the `xsengine.ini` file. If no such key exists, the default session timeout of 900 seconds is used. If you want to define a non-default value for the scheduler's `sessiontimeout` key, you must create the key in the `scheduler` section of the `xsengine.ini` file and supply the desired timeout value, for example, 600 seconds.

Caution

It is not recommended to specify a value of 0 (zero) for the `sessiontimeout` key; this disables the session-timeout feature for all jobs started by the scheduler.

c. Save the job.

Choose *Save Job* to save and activate the changes to the job schedule.

6. Enable the job-scheduling feature in SAP HANA XS.

This step requires the permissions granted to the SAP HANA administrator.

i Note

It is not possible to enable the scheduler for more than one host in a distributed SAP HANA XS landscape.

- a. In the *XS Job Dashboard* set the *Scheduler Enabled* toggle button to **YES**.

Toggling the setting for the *Scheduler Enabled* button in the *XS Job Dashboard* changes the value set for the SAP HANA configuration variable `xsengine.ini > scheduler > enabled`, which is set in the *Configuration* tab of the SAP HANA studio's *Administration* perspective.

7. Check the job logs to ensure the XS job is active and running according to the defined schedule.

You can view the `xsjob` logs in the *XS Job Dashboard* tab of the *SAP HANA XS Administration Tool*.

i Note

To maintain details of an XS job using the Web-based *XS Administration Tool* you need the privileges granted in the SAP HANA user role `sap.hana.xs.admin.roles::JobAdministrator`.

If the job does not run at the expected schedule, the information displayed in the `xsjob` logs includes details of the error that caused the job to fail.

Related Information

[The XS Job File \[page 1630\]](#)

12.1.11.3.1 The XS Job File

The `.xsjob` file defines the details of a task that you want to run (for example, an XS JavaScript or an SQLScript) at a scheduled interval.

The XS job file uses a `cron`-like syntax to define the schedule at which the service defined in an XS JavaScript or SQLScript must run, as you can see in the following example, which runs the specified job (the stock-price checking service `yahoo.xsjs`) on the 59th second minute of every minute.

```
{
  "description": "Read stock value",
  "action": "yahoo:yahoo.xsjs::readStock",
  "schedules": [
    {
      "description": "Read current stock value",
      "xscron": "* * * * * 59",
      "parameter": {
        "stock": "SAP.DE"
      }
    }
  ]
}
```

When defining the job schedule in the `xsjob` file, pay particular attention to the entries for the following keywords:

- `action`
Text string used to specify the path to the function to be called as part of the job.

```
"action": "<package_path>:<XSJS_Service>.xsjs::<FunctionName>",
```

Note

You can also call SQLScripts using the `action` keyword.

- `description`
Text string used to provide context when the XSjob file is displayed in the *SAP HANA XS Administration* tool.
- `xscron`
The schedule for the specified task (defined in the "action" keyword); the schedule is defined using cron-like syntax.
- `parameter`
A value to be used during the action operation. In this example, the parameter is the name of the stock `SAP.DE` provided as an input for the parameter (`stock`) defined in the `readStock` function triggered by the `xsjob` action. You can add as many parameters as you like as long as they are mapped to a parameter in the function itself.

The following examples illustrate how to define an `xscron` entry including how to use expressions in the various `xscron` entries (day, month, hour, minutes,...):

- `2013 * * fri 12 0 0`
Every Friday of 2013 at 12:00 hours
- `* * 3:-2 * 12:14 0 0`
Every hour between 12:00 and 14:00 hours on every day of the month between the third day of the month and the second-last day.

Tip

In the day field, third from the left, you can use a negative value to count days backwards from the end of the month. For example, `* * -3 * 9 0 0` means: three days from the end of every month at 09:00.

- `* * * * * */5 *`
Every five minutes (`*/5`) and at any point (`*`) within the specified minute.

Note

Using the asterisk (`*`) as a wild card in the seconds field can lead to some unexpected consequences, if the scheduled job takes less than 59 seconds to complete; namely, the scheduled job restarts on completion. If the scheduled job is very short (for example, 10 seconds long), it restarts repeatedly until the specified minute ends.

To prevent short-running jobs from restarting on completion, schedule the job to start at a specific second in the minute. For example, `* * * * * */5 20` indicates that the scheduled job should run every five minutes and, in addition, at the 20th second in the specified minute.

- `* * * -1.sun 9 0 0`
Every last Sunday of a month at 09:00 hours

Related Information

[Tutorial: Schedule an XS Job \[page 1627\]](#)

12.1.12 Maintaining Translation Text Strings

Maintain the translated text strings used in an application's user interface, error messages, and documentation.

For the purposes of localisation (L10N), you can provide the text strings displayed in an application's user interface in multiple languages, for example, English, French, or Chinese. You can also provide notifications and error messages in the same, local languages. To manage and maintain these translated text strings, SAP HANA provides an online translation tool (OTT). The translation of the text strings themselves can be performed manually or with suggestions provided by an external service, for example, SAP Translation Hub. Access to external translation services is not covered by the SAP HANA license and usually requires a user account.

Setting up and maintaining the online translation tools for SAP HANA includes the following high-level tasks:

- Enabling the translation tool
- Accessing packages in the SAP HANA repository
- Maintaining text strings in the source and target languages
This tasks involves maintaining the contents of the following SAP HANA tables:
 - ACTIVE_CONTENT_TEXT
 - ACTIVE_CONTEXT_TEXT_CONTENT
 - ACTIVE_OBJECT_TEXT
 - ACTIVE_OBJECT_TEXT_CONTENT
- Enabling access to a remote text-translation service (**optional**)

! Restriction

Access to external translation services is not granted in the SAP HANA license. To use external translation services such as the *SAP Translation Hub*, an additional license is required. In addition, the *SAP Translation Hub* is currently available only for Beta testing.

- Maintaining HTTP destinations for any remote systems that provide services used by the *Online Translation Tool* (**optional**)
Remote translation services such as *SAP Translation Hub* can provide access to a database of translated text strings, which are used to provide suggestions in the target language. To access such a remote service, you must maintain an HTTP destination (or extend an existing destination) that provides details of the host system where the translation service is running as well as a valid user account and logon authentication. You must also ensure that a trust relationship exists between the translation server and SAP HANA, for example, by importing the translation server's client certificate into the SAP HANA trust store.

The SAP HANA *Online Translation Tool* is available on the SAP HANA XS Web server at the following URL:

`http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/translationTool/`

→ Tip

The privileges required to use the SAP HANA *Online Translation Tool* (OTT) are granted by the role `templatesap.hana.xs.translationTool.roles::translator`.

Related Information

[Create and Edit Text Translations \[page 1633\]](#)

[Export and Import Translated Text \[page 1637\]](#)

[SAP Translation Hub Cloud Service \(beta\)](#)

12.1.12.1 Create and Edit Text Translations

Maintain translations for text strings displayed in an SAP HANA application's user interface.

Prerequisites

To maintain translated text for an application in SAP HANA XS, the following prerequisites apply:

- You have access to an SAP HANA system.
- You have the privileges required to access the repository packages containing the text strings to be localized/translated.
- You have a role based on the role template `sap.hana.xs.translationTool.roles::translator`.
- If you want to make use of optional external translation services, you must maintain access to the translation server system.

! Restriction

Access to external translation services is not granted in the SAP HANA license. To use external translation services such as the *SAP Translation Hub*, an additional license is required. The *SAP Translation Hub* is currently available only for BETA testing.

Details of the remote systems where the translation service is running (for example, SAP Translation Hub) are defined in HTTP destination configuration files along with details of any corresponding user account and authentication certificates.

Context

An application's user interface and notifications can be translated from the original source language (for example, English) into one or more local (target) languages, for example, French, Spanish, or Japanese. You

can either translate the texts manually or with the help of an (optional) external translation service. To provide translations of the UI text strings for your SAP HANA application, perform the following steps:

Procedure

1. Start the *SAP HANA Online Translation Tool*.

The *SAP HANA Online Translation Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/translationTool`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must also have the privileges required to perform the tasks associated with the maintenance of translation texts.

2. Select the delivery unit that contains the application with the text strings you want to translate.

Use the *Delivery Unit* drop-down list to select a delivery unit.

→ Tip

The name of the vendor associated with the selected delivery unit is displayed automatically in the *Vendor* field, for example, *acme.com*; the vendor name cannot be changed here.

3. Select the package that contains the text strings you want to translate.

Use the *Package* drop-down list to select a package. If the selected package contains text elements, they are displayed alphabetically in a list.

→ Tip

The original source language associated with the contents of the selected package is displayed automatically.

4. Enable access to a text-translation service, for example, *SAP Translation Hub*. (**optional**).

! Restriction

Access to external translation services is not granted in the SAP HANA license. To use external translation services, an additional license is required.

If you want to make use of the services provided by a translation server, you need to maintain an HTTP destination **extension** that provide details of the host system where the translation service is running; access to the translation service usually requires a user account and logon authentication. You must also ensure that a trust relationship exists between the translation server and SAP HANA, for example, by importing the translation server's client certificate into the SAP HANA trust store that you are using to handle authentication for this HTTP destination.

The HTTP destination configuration

`sap.hana.xs.translationTool.server:translationService.xshttpdest` defines details of the server hosting the SAP Translation Hub service. Although you cannot edit this destination configuration,

note that you can use an HTTP destination **extension** to change the details, for example, to point to an alternative host name.

5. Add a translation for a text element.

For a given text element in the *Text ID* list, you can provide a suitable translation in one or more languages, for example: French (*fr*), Spanish (*es*), and Japanese (*ja*).

- a. Expand the desired UI text element.

In the *Text ID* list, locate and expand the element for which you want to provide a translation.

- b. Add a translation.

Choose *Add Translation*.

- c. Select the desired language for the translation from the *Target Language* drop-down list.

- d. In the *Target Language Text* box, type the translation for the selected text element.

→ Tip

If the *SAP Translation Hub* option is enabled, language-specific suggestions for possible translation matches are provided as you type. If you see a suggestion that is suitable, use the mouse to select the suggested text.

- e. Add another translation.

Choose *Add Translation*

- f. Edit an existing translation

Choose the *Edit* icon next to the translation you want to modify and make the required changes.

6. Save your additions and changes.

Choose *Save* to store the added translations or any modifications in the appropriate tables in the SAP HANA database.

Related Information

[Online Translation Tool Details \[page 1635\]](#)

[Export and Import Translated Text \[page 1637\]](#)

[Edit an HTTP Destination Runtime Configuration \[page 1546\]](#)

[Managing Trust Relationships \[page 1552\]](#)

12.1.12.1.1 Online Translation Tool Details

Display details of the source text for an application's user interface elements and, if available, any available translations.

The *Online Translation Tool* tool enables you to view details of the text elements contained in the individual packages of an SAP HANA application. The following table indicates which information can be viewed.

i Note

The privileges required to use the SAP HANA *Online Translation Tool* (OTT) are granted by the role template `sap.hana.xs.ott.roles::translator`.

Translation Text Details

UI Element	Description	Example
<i>Delivery Unit</i>	Name of the SAP HANA delivery unit (DU) that contains the default text strings for which a translation is required along with the name of the vendor associated with the selected delivery unit	ACME_XS_BASE - acme.com
<i>Package</i>	The name of (and path to) the package containing the text strings for which a translation is required	acme.com.app.ui.login
<i>Source language</i>	Short name of the source language for the text strings contained in the selected package, for example: en (English), fr (French), ja, (Japanese)	en
<i>Target Language</i>	Long or short name of the target language for the text strings contained in the selected package, for example: Bulgarian (bg), French (fr), Japanese (ja)	Chinese (zh)
<i>Domains</i>	The SAP product-specific translation domain to which the selected DU/package belongs, for example, <i>Financial Accounting</i> or <i>Customer Relationship Managment</i> . Domains are used in the translation process to determine the correct terminology for a text string that has to be translated; the same text might require a different translation depending on the domain (or application) in which it is used. Suggestions from a remote translation service such as the SAP Translation Hub are restricted to the currently selected domain.	"Basis", or "Accounting - General"
<i>Enable Translation Hub</i>	Enable automatic suggestions (in the <i>Target language text</i> box) for translation texts using a remote service such as SAP Translation Hub; the suggestions are provided by a remote translation database.	Yes/No

! Restriction

Access to external translation services is not granted in the SAP HANA license. To use external translation services such as the *SAP Translation Hub*, an additional license is required. The *SAP Translation Hub* is currently available only for BETA testing.

Access to the remote translation service usually requires a user account and logon authentication. You also need to maintain an HTTP destination (or extend an existing one) for the translation server system and ensure the server system is trusted by SAP HANA, for example, by importing the translation server's client certificate into the SAP HANA trust store.

UI Element	Description	Example
<i>Text ID</i>	The name/ID of the UI element for which a text string is required. This could be a tab title, a box name, a notification, or an error message.	LOGON_LABEL
<i>Default Text</i>	The text string associated with the text ID	HANA Logon
<i>Target Language Text</i>	Proposed/accepted translation (in the target language) of the text string displayed (in the source language) in the <i>Default Text</i> field. Activate the <i>Enable Translation Hub</i> option to enable auto-suggestions in the target language.	-
<i>Source Object</i>	The name of the design-time artifact that contains the UI text strings.	logonForm.hdbtextbundle

Related Information

[Create and Edit Text Translations \[page 1633\]](#)

[Export and Import Translated Text \[page 1637\]](#)

[Managing Trust Relationships \[page 1552\]](#)

12.1.12.2 Export and Import Translated Text

Transport text translations between systems using the industry-standard, XML-based `xliff` format.

Prerequisites

To export and import translated text for an application in SAP HANA XS, the following prerequisites apply:

- You have access to an SAP HANA system.
- You have access to the repository packages containing the text strings to be localized/translated.
- You have been granted a role based on the role template `sap.hana.xs.translationTool.roles::translator`.

Context

An application's user interface and notifications can be translated from the original source language (for example, English) into one or more target local languages, for example, French, Spanish, or Japanese. To provide translations of the UI text strings for your SAP HANA application, perform the following steps:

Procedure

1. Start the *SAP HANA Online Translation Tool*.

The *SAP HANA Online Translation Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/translationTool`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must also have the privileges required to perform the tasks associated with the maintenance of translation texts.

2. Select the delivery unit that contains the application with the text strings you want to translate.

Use the *Delivery Unit* drop-down list to select a delivery unit.

→ Tip

The name of the vendor associated with the selected delivery unit is displayed automatically in the *Vendor* field, for example, *acme.com*. You cannot change this here.

3. Select the package that contains the text strings you want to translate.

Use the *Package* drop-down list to select a package. If the selected package contains text elements, they are displayed automatically in an alphabetically ordered list.

→ Tip

The original source language associated with the contents of the selected package is displayed automatically.

4. Export the UI text elements from the local source system.

You can export the translation texts to an archive on a local file system using the industry-standard, XML-based `xliff` format.

5. Import the UI text elements to the remote target system.

You can import the translation texts into SAP HANA from an archive whose content are stored using the industry-standard, XML-based `xliff` format.

6. Confirm that the import operation was successful.

Check the status of the following tables in the SAP HANA database:

- ACTIVE_CONTENT_TEXT
- ACTIVE_CONTEXT_TEXT_CONTENT

- ACTIVE_OBJECT_TEXT
- ACTIVE_OBJECT_TEXT_CONTENT

Related Information

[Online Translation Tool Details \[page 1635\]](#)

[Create and Edit Text Translations \[page 1633\]](#)

12.1.13 Maintaining HTTP Traces for SAP HANA XS Applications

HTTP tracing for individual SAP HANA XS applications can be enabled on the SAP HANA Web Dispatcher.

The *SAP HANA XS Administration Tools* include the *SAP Web Dispatcher HTTP Tracing* application, which you can use to enable and disable HTTP tracing on the SAP Web Dispatcher for SAP HANA XS applications.

Note

SAP HANA uses roles to grant access to the features provided by the *SAP HANA XS Administration Tool*. To access the administration tools required to manage HTTP tracing on the SAP Web Dispatcher, you must have a role based on the role template `WebDispatcherHTTPTracingAdministrator`. The role template `WebDispatcherHTTPTracingViewer` contains the privileges for read-only access to the *SAP Web Dispatcher HTTP Tracing* tool.

You can use the *SAP HANA XS Administration Tools* to perform the following tasks:

- Display a list of all traced applications
List all applications defined in the system. Details include the application's metadata, information about HTTP tracing configuration for the particular application, the status of the XS job that starts the tracing process, and HTTP tracing log information.
- Enable HTTP tracing
Enable HTTP tracing for selected SAP HANA XS applications
- Disable HTTP tracing
Disable HTTP tracing for selected SAP HANA XS applications

Tracing is managed by the XS job `sap.hana.xs.admin.webdispatcher.jobs::httptracing.xsjob`, which runs at a predefined schedule. If you enable or disable HTTP tracing, you must modify the XS job file accordingly.

→ Tip

Administrator access to the XS job details requires the privileges granted by the role template `sap.hana.xs.admin.roles::JobAdministrator`. These privileges are already included in the `WebDispatcherHTTPTracingAdministrator` role template, which is required to use the *SAP Web Dispatcher HTTP Tracing*.

HTTP tracing is enabled by setting configuration parameters in SAP HANA XS (`xsengine.ini`) and the SAP Web Dispatcher (`webdispatcher.ini`). If an SAP HANA XS application is defined in a parameter in

`xsengine.ini`, then HTTP tracing is enabled for the specified application. If not, then HTTP tracing is disabled for the application.

i Note

If HTTP tracing is disabled for an application, the corresponding HTTP trace parameters in `xsengine.ini` and `webdispatcher.ini` are removed. If you re-enable HTTP tracing on the SAP Web Dispatcher for the same application, the required parameters are recreated automatically.

Connections to the database are performed with the SQL auto-user defined in `/sap/hana/xs/admin/webdispatcher/server/common/httpTracing.xssqlcc`.

Related Information

[SAP HANA XS Classic Administration Roles \[page 1529\]](#)

[Enable HTTP Tracing for an SAP HANA XS Application \[page 1642\]](#)

[Maintain XS Job Details \[page 1619\]](#)

12.1.13.1 Display the HTTP Trace Status of SAP HANA XS Applications

Display a list of SAP HANA XS applications which shows the status of HTTP tracing.

Prerequisites

To use the *SAP HANA XS Administration Tool* to view the current status of HTTP tracing for SAP HANA XS applications, the following prerequisites apply:

- You have administrator access to an SAP HANA system.
- You have been granted roles based on one of the following role templates:
 - `sap.hana.xs.admin.roles::WebDispatcherHTTPTracingViewer`
 - `sap.hana.xs.admin.roles::WebDispatcherHTTPTracingAdministrator`

Context

To use the *SAP HANA XS Administration Tool* to display a list of applications and the HTTP trace status, perform the following steps:

Procedure

1. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must have the privileges required to perform administration tasks with the *SAP Web Dispatcher HTTP Tracing* tool.

2. Start the *SAP Web Dispatcher HTTP Tracing* tool.

In the list of XS Administration tools, choose *SAP Web Dispatcher HTTP Tracing*.

3. Display a list of the SAP HANA XS applications running on the system to which you are connected; the *HTTP Tracing Enabled* column indicates (Yes/No if HTTP tracing is enabled for the application).

→ Tip

You can use the search box to display a list of only those applications that match a particular string, for example, "**admin**".

Related Information

[Enable HTTP Tracing for an SAP HANA XS Application \[page 1642\]](#)

[Application HTTP Tracing Details \[page 1641\]](#)

12.1.13.1.1 Application HTTP Tracing Details

Display a list of the SAP HANA XS applications for which HTTP tracing is enabled on the SAP Web Dispatcher.

The *XS Applications* tab in the *SAP Web Dispatcher HTTP Tracing* tool enables you to view a list of the SAP HANA XS applications for which HTTP tracing is enabled on the SAP Web Dispatcher. The following table indicates which information can be viewed.

→ Tip

You can use the search box to display a list of the applications that match a particular string, for example, “**admin**”.

Job Details

UI Element	Description	Example
<i>SAP Web Dispatcher HTTP Tracing Job</i>	The SAP HANA XS job used to start the tracing operation for the listed applications	httptracing.xsjob
<i>ACTIVE/INACTIVE</i>	The current status of the HTTP tracing job that manages the tracing operation for the selected applications	ACTIVE
<i>Application Name</i>	The full path to (and the name of) the SAP HANA XS application for which HTTP tracing is enabled on the SAP Web Dispatcher	sap.hana.xs.admin
<i>Delivery Unit</i>	The name of the delivery unit that contains the application specified in <i>Application Name</i>	HANA_XS_ADMIN
<i>Vendor</i>	The name of the vendor responsible for the creation and maintenance of the delivery unit that contains the traced application	sap.com
<i>HTTP Tracing Enabled</i>	The current tracing status: No (disabled); yes (enabled)	Yes

Related Information

[Enable HTTP Tracing for an SAP HANA XS Application \[page 1642\]](#)

12.1.13.2 Enable HTTP Tracing for an SAP HANA XS Application

HTTP tracing on the SAP Web Dispatcher can be enabled for one or more SAP HANA XS applications

Prerequisites

To enable HTTP tracing on the SAP Web Dispatcher for SAP HANA XS applications, the following prerequisites apply:

- You have administrator access to an SAP HANA system.
- You have been granted a role based on the role template `sap.hana.xs.admin.roles::WebDispatcherHTTPTracingAdministrator`.
- The XS job `sap.hana.xs.admin.webdispatcher.jobs::httptracing.xsjob` is configured and running. (By default, the job runs at 12:00 every day.)

- The XS SQL connection configuration `/sap/hana/xs/admin/webdispatcher/server/common/httpTracing.xssqlcc` is active (available by default).

Context

To enable HTTP tracing on the SAP Web Dispatcher for an application, you must perform the following steps:

Procedure

1. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must have the privileges required to perform administration tasks with the *SAP Web Dispatcher HTTP Tracing* tool.

2. Start the *SAP Web Dispatcher HTTP Tracing* tool.
In the list of XS administration tools, choose *SAP Web Dispatcher HTTP Tracing*.
3. Display a list of the SAP HANA XS applications running on the system to which you are connected.

Tip

You can use the search box to display a list of only those applications that match a particular string, for example, `admin`.

4. Enable HTTP tracing for an application.

In the *XS Applications* tab, the *HTTP Tracing Enabled* column indicates if HTTP tracing is enabled or not (*Yes/No*) for the application.

- a. In the *XS Applications* tab, choose *Edit*.
- b. **Check** the box for the application for which you want to enable HTTP tracing.
- c. In the *XS Applications* tab, choose *Save*.

Saving the changes to the configuration enables HTTP tracing and automatically sets the following configuration parameters (keys):

- Configuration section: `webdispatcher.ini/profile`
 - key
`icm/HTTP/logging_n`
 - value
`PREFIX=/path/to/app/, LOGFILE=$(_LOCAL_HOST_NAME)/trace/access_log_app-%y-%m-%d, MAXSIZEKB=10000, SWITCHTF=day, LOGFORMAT=SAP, FLUSH=1`

- Configuration section: `xsengine.ini/customer_usage`
 - `key=/path/to/appname/`
 - `value=icm/HTTP/logging_n`

→ Tip

This is the value defined for the key `webdispatcher.ini/profile`.

5. Update the XS job used to start the trace operation.

The XS job `sap.hana.xs.admin.webdispatcher.jobs:httptracing.xsjob` is used to stop and start HTTP tracing on the SAP Web Dispatcher for individual XS applications. The current status of the XS job is indicated in the *SAP Web Dispatcher HTTP Tracing* dialog.

- a. In the *SAP Web Dispatcher HTTP Tracing* dialog, click the link to the XS job `sap.hana.xs.admin.webdispatcher.jobs:httptracing.xsjob`.

The *XS Job Details* window displays a brief description of the XS job and information about any configured schedules.

- b. Choose the *Configuration* tab to set up the XS job.
- c. Type the name of a user with the required permission to run the XS job and the corresponding password.
- d. Check the *Active* box.
- e. Choose *Save Job* to update the XS job and start the HTTP tracing.

i Note

A user name and password are required to save the changes you make to the XS job.

6. Check the new log file is created and contains entries.

The log file is located in the folder you specified in the `webdispatcher.ini/profile` key `icm/HTTP/logging_n`, for example:

```
LOGFILE=$( _LOCAL_HOST_NAME )/trace/access_log_app-%y-%m-%d
```

Where `app` is the name of the application whose HTTP traffic you are tracing.

Related Information

[Application HTTP Tracing Details \[page 1641\]](#)

[SAP HANA XS Classic Configuration Parameters \[page 1532\]](#)

[SAP HANA XS Classic Administration Roles \[page 1529\]](#)

12.1.13.3 Disable HTTP Tracing for an SAP HANA XS Application

HTTP tracing on the SAP Web Dispatcher can be disabled for one or more SAP HANA XS applications.

Prerequisites

To enable HTTP tracing on the SAP Web Dispatcher for SAP HANA XS applications, the following prerequisites apply:

- You have administrator access to an SAP HANA system.
- You have been granted a role based on the role template `sap.hana.xs.admin.roles::WebDispatcherHTTPTracingAdministrator`.
- The XS job `sap.hana.xs.admin.webdispatcher.jobs::httptracing.xsjob` is configured and running. (By default, the job runs at 12:00 every day.)
- The XS SQL connection configuration `/sap/hana/xs/admin/webdispatcher/server/common/httpTracing.xssqlcc` is active (available by default).

Context

To disable HTTP tracing on the SAP Web Dispatcher for an application, you must perform the following steps:

Procedure

1. Start the *SAP HANA XS Administration Tool*.

The *SAP HANA XS Administration Tool* tool is available on the SAP HANA XS Web server at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/xs/admin/`.

i Note

In the default configuration, the URL redirects the request to a logon screen, which requires the credentials of an authenticated SAP HANA database user to complete the logon process. The user who logs on must have the privileges required to perform administration tasks with the *SAP Web Dispatcher HTTP Tracing* tool.

2. Start the *SAP Web Dispatcher HTTP Tracing* tool.
In the list of XS administration tools, choose *SAP Web Dispatcher HTTP Tracing*.
3. Display a list of the SAP HANA XS applications running on the system to which you are connected.

→ Tip

You can use the search box to display a list of only those applications that match a particular string, for example, "`admin`".

4. Disable HTTP tracing for an application.

In the *XS Applications* tab, the *HTTP Tracing Enabled* column indicates if HTTP tracing is enabled or not (*Yes/No*) for the application.

- a. In the *XS Applications* tab, choose *Edit*.
- b. **Uncheck** the box for the application for which you want to **disable** HTTP tracing.
- c. In the *XS Applications* tab, choose *Save*.

Saving the changes to the configuration disables HTTP tracing for the selected application and **removes** the following parameters (keys):

- Configuration section: `webdispatcher.ini/profile`
 - `key=icm/HTTP/logging_n`
- Configuration section: `xsengine.ini/customer_usage`
For example:
 - `key=/path/to/appname/`

5. Update the XS job used to stop the trace operation.

The XS job `sap.hana.xs.admin.webdispatcher.jobs::httptracing.xsjob` is used to stop and start HTTP tracing on the SAP Web Dispatcher for individual XS applications. The current status of the XS job is indicated in the *SAP Web Dispatcher HTTP Tracing* dialog.

- a. In the *SAP Web Dispatcher HTTP Tracing* dialog, click the link to the XS job `sap.hana.xs.admin.webdispatcher.jobs::httptracing.xsjob`.

The *XS Job Details* window displays a brief description of the XS job and information about any configured schedules.

- b. Choose the *Configuration* tab to set up the XS job.
- c. Type the name of a user with the required permission to run the XS job and the corresponding password
- d. Uncheck the *Active* box.
- e. Choose *Save Job* to update the XS job and stop the HTTP tracing.

i Note

A user name and password are required to save the changes you make to the XS job.

6. Check that tracing has been switched off and **no** new logs files are being created.

The log files for the traced application are located in the folder you specified in the `webdispatcher.ini/profile` key `icm/HTTP/logging_n`, for example:

```
LOGFILE=$( _LOCAL_HOST_NAME )/trace/access_log_app-%y-%m-%d
```

Where `app` is the name of the application whose HTTP traffic you are tracing.

Related Information

[SAP HANA XS Classic Administration Roles \[page 1529\]](#)

[SAP HANA XS Classic Configuration Parameters \[page 1532\]](#)

12.2 Maintaining the SAP HANA XS Advanced Model Run Time

Maintain the SAP HANA XS advanced model run-time environment.

From HANA 1.0 SPS 11, SAP HANA includes an additional run-time environment for application development: SAP HANA extended application services (XS), advanced model. SAP HANA XS, advanced model represents an evolution of the application server architecture within SAP HANA by building upon the strengths (and expanding the scope) of SAP HANA extended application services (XS), classic model. SAP recommends that customers and partners who want to develop new applications use SAP HANA XS advanced model.

→ Tip

If you want to migrate existing XS classic applications to run in the new XS advanced run-time environment, SAP recommends that you first check the features available with the installed version of XS advanced; if the XS advanced features match the requirements of the XS classic application you want to migrate, then you can start the migration process.

SAP HANA extended application services, advanced model (XS advanced) is a platform that enables the management of polyglot Web-based applications built on micro-services but also supports the more traditional "monolithic" application. That means that, while XS classic applications were restricted to a server-side JavaScript dialect called XSJS, applications in XS advanced can be written in a programming language of your choice. By default, the platform supports Java, Node.js and XSJS, and Python, but also contains a plug-in concept that enables the use of additional "custom" languages, too. Moreover, in contrast to XS classic, XS advanced applications are not executed by a central server. Instead, each application has its own server and run-time environment – an architecture that enables a very flexible composition of applications using the concept of micro-services.

In XS advanced, an application can either be comprised of several independent and replaceable components or it can reuse services that are already present on the platform, and each component can be developed, maintained, and updated individually. This allows a high level of flexibility in terms of deployment. For example, an application can be composed of a Java and a Node.js component, each of which can be developed and updated independently, and, if required, the same application can reuse an additional service written in Python.

It's the task of the XS advanced run time to try to reduce the additional complexity this advanced concept might introduce to a minimum. For example, the XS run-time environment is responsible for connecting micro-services together to form a complete application. XS advanced starts, stops, and monitors applications, and provides a simple means of connecting applications to backing services such as the SAP HANA database (using a concept called "service brokers", which provide components for user authentication and the higher level application deployment services). Note that XS advanced is designed in a way that makes it possible to write business applications that can be moved from the on-premise platform to the SAP Cloud Platform and vice versa.

→ Tip

For general information, advice, and a list of frequently asked questions about XS advanced issues and solutions, see *Related Information* below.

Downloading XS Advanced from SAP Marketplace

SAP HANA Extended Application Services, advanced model, is available not only on the SAP HANA media but also as a separate component on SAP Marketplace. Users with the required S-User ID can download the latest version of XS advanced component in the package `SAP_EXTENDED_APP_SERVICES_1` from the following location:

▶ [Service Marketplace](#) ▶ [Software Downloads \[Downloads\]](#) ▶ [SUPPORT PACKAGES & PATCHES](#) ▶ [By Alphabetical Index \(A-Z\)](#) ▶ [H](#) ▶ [SAP HANA PLATFORM EDITION](#) ▶

- ▶ [SAP HANA PLATFORM EDITION 1.0](#) ▶ [XS ADVANCED RUNTIME](#) ▶ [SAP EXTENDED APP SERVICES 1](#) ▶
- ▶ [SAP HANA PLATFORM EDITION 2.0](#) ▶ [SAP EXTENDED APP SERVICES 1](#) ▶

→ Tip

SAP HANA Extended Application Services, advanced model, is backwards compatible; you can provide access to new features by installing the latest version of the XS advanced component even on older versions of SAP HANA. To download the package `SAP_EXTENDED_APP_SERVICES_1`, see *SAP Software Download Center* in *Related Information* below.

Related Information

[SAP Note 2596466 - FAQ: SAP HANA XS Advanced](#)

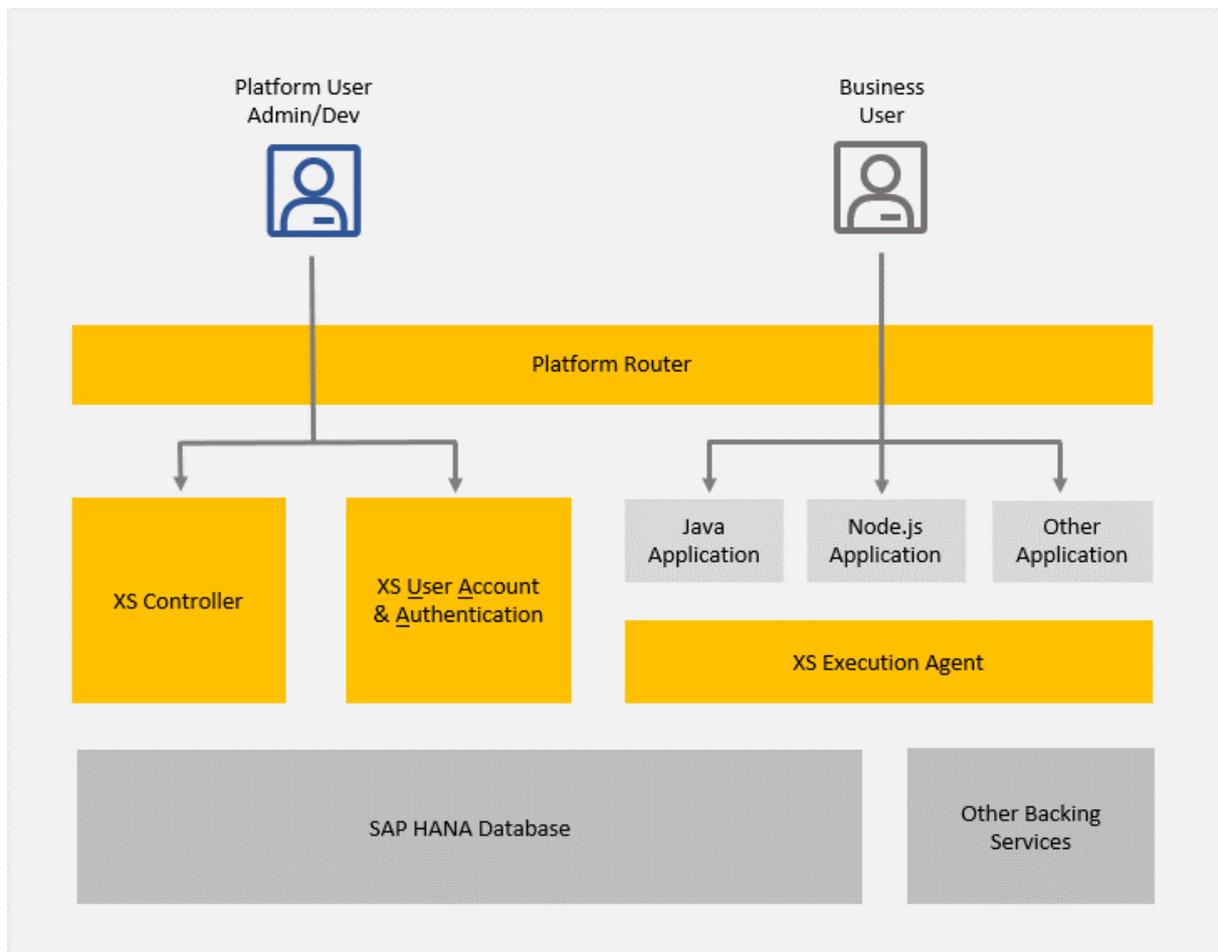
[SAP Software Download Center \(Logon required\)](#)

[2542036 - XS Advanced SPS 03 Release Note](#)

12.2.1 XS Advanced Platform Components

An overview of the main components of the XS advanced model run-time environment.

The SAP HANA XS advanced model run-time environment is comprised of the components illustrated in the following diagram and described in more detail in the corresponding sections.



The XS Advanced Controller

The XS advanced Controller (XS Controller) is a central component of the platform and is deployed once within an SAP HANA system. The XS Controller has the following purposes:

1. The XS Controller provides an application programming interface (API) for the platform user (administrators or application developers). The API enables users to perform the following tasks:
 - Maintain applications and service instances (resources provided by backing services)
 - Maintain XS advanced role collections (by means of an API provided by the XS User Account and Authentication service)
 - Maintain application run-time environments
 - Maintain application build packs (a convenient way to bring in custom programming languages for applications)

→ Tip

A command-line based administration tool (the `xs` client) and a GUI based tool (the XS Advanced Admin Cockpit) are provided to help administrators maintain the XS advanced model run-time platform. For more information, see *Related Information* below.

2. The XS Controller maintains the state of the XS advanced platform and communicates with a group of Execution Agents which are used to start, monitor, or stop applications running in the XS advanced run-time environment.

The XS Advanced Execution Agent

The Execution Agent manages and monitors application instances. Several Execution Agents can be connected to the XS Controller, for example, if it is necessary to scale out application instances to several hosts.

Managed by the `xsexecagent` service, Execution Agents are primarily responsible for starting and stopping application instances. To be reachable for end users, launched application instances typically provide a public HTTP port. The Execution Agent checks the availability of this end point regularly. Instances of an application that are no longer reachable are restarted automatically.

In a distributed system, different instances of the same application do not necessarily run on the same host, although it is possible to pin applications to a particular host. If spaces are mapped to different OS users, the Execution Agents also ensure that application instances running in different spaces are not visible to each other at the operating system (OS) level.

XS Advanced User Account and Authentication

The XS advanced User Account and Authentication (UAA) is the component that handles user-authentication requests for platform users as well as for business users.

The UAA service is a multi-tenant identity management service. The UAA's primary role is as an OAuth2 provider, issuing tokens for client applications to use when they act on behalf of XS advanced (or Cloud Foundry) users. The UAA can also authenticate users with their XS advanced (or Cloud Foundry) credentials and can act as an SSO service using those credentials (or others).

As well as various other management functions, the UAA provides end points for the management of user accounts and for registering OAuth2 clients.

You can configure an SAP HANA system to act as a service provider for XS advanced applications that use Single Sign On (SSO) authentication based on Security Assertion Markup Language (SAML) certificates. An SAML identity provider is used by the SAML service provider to authenticate users signing in by means of a single sign-on mechanism. If SAP HANA XS advanced applications are configured to use SAML assertions as the logon authentication method, then an SAML IdP is required. It is also possible to assign roles automatically to users who log on to an application by means of single sign-on with SAML assertions.

→ Tip

For more information about maintaining SAML Identity Providers (IdP) in XS advanced, see *Related Information* below.

XS Advanced Platform Router

The Platform Router is the Web-dispatcher component that exposes all public end points of XS advanced platform components as well as for the running applications. The Platform Router is also responsible for load-balancing requests between several instances of an application. One or more end points (so called routes) can be exposed per application.

XS Advanced Application Instances

By default, the XS advanced platform supports applications written in Java, Node.js, server-side JavaScript (XSJS), and Python, as well as other “custom” languages.

When an application is started by a platform user, the Execution Agent starts one or more instances of this application. Each application is run in its own environment and file-system sandbox. In this way, several application instances can be started, for example, depending on the current load an application has to take, without affecting the other application instances.

The number of application instances can be adapted (scaled) during run time. All instances of an applications are exposed under the same route, which means they share the same URI, a feature that enables transparent load-balancing. Using a round-robin algorithm, the Platform Router balances the load by distributing requests between available application instances. Sticky sessions are used to ensure proper caching of data inside single application instances.

XS Advanced Backing Services

Backing Services provide resources to applications. Arbitrary backing services can be connected to the XS advanced platform by so-called "Service Brokers". The service broker provides a standardized interface to the XS Controller to fetch credentials for a particular backing service. The credentials are exposed to applications by means of so-called "service instances".

→ Tip

For more information about the XS advanced Service Broker, see *Related Information* below.

A platform user can create a service instance and bind this service instance to a particular application, while in the background the XS Controller will fetch the corresponding credentials. When starting the application, the XS Controller injects these credentials into the environment of the application instance, which enables the application to use the backing service.

SAP HANA Database

The XS advanced platform uses SAP HANA as persistence as well as an infrastructure provider to manage the life cycle of XS advanced services, so it is important that you monitor the operation of SAP HANA databases on a regular basis.

Although SAP HANA actively alerts you about critical problems and situations, keeping an eye on resource usage and performance can help you identify patterns, forecast requirements, and recognize when something is wrong.

You can monitor SAP HANA using the SAP HANA Cockpit. The tools in the SAP HANA Cockpit rely on the monitoring and alerting information provided by system views and the statistics service.

Platform Administration

During installation of the SAP HANA XS advanced run-time component, a user named `XSA_ADMIN` is created and granted the privileges required to administrate the XS advanced run-time environment. The `XSA_ADMIN` user can be used to log in to the XS advanced platform, perform administration tasks, or create further administration or restricted users, for example, using the following tools:

- Command-line tools:
 - `xs`
Maintain XS advanced run-time components (applications, services, brokers, etc.)
 - `XSA`
Stop, start XS advanced run-time instances.

i Note

Some `XSA` command require `XSA_ADMIN` user privileges; most `XSA` commands require only `<sid>adm` privileges, for example, `XSA restart`.

- GUI-based tools:
XS Advanced Admin Cockpit.

Related Information

[The XSA Command Reference \[page 1740\]](#)

[The XS Command-Line Interface \[page 1653\]](#)

[Maintaining the XS Advanced Runtime Environment with SAP HANA XS Advanced Cockpit \[page 1753\]](#)

[Managing SAML Identity Providers in XS Advanced \[page 1782\]](#)

12.2.2 Maintaining the XS Advanced Run-time Environment with a Command-Line Interface

Use command-line tools to administrate and maintain XS advanced-model run-time components.

SAP HANA provides the following tools to enable you to maintain your XS advanced run-time environment from the command line:

- The `xs` command-line interface

A selection of utilities to help you maintain not only the applications that are deployed to the XS advanced run-time environment, but also the run-time environment itself, for example, the organizations and spaces, and the users who need access and use it.

```
xs <command> [<ARGUMENTS>] [<OPTIONS>]
```

The `xs` CLI is automatically installed along with the XS advanced platform on an SAP HANA system. If the `xs` CLI is not added to the SAP HANA system's `<PATH>` environment variable, you can find the executable in the default location `/bin/xs`. The `xs` CLI can also be used from a remote computer which connects over a secure connection to the SAP HANA system where you want to perform the administration tasks.

→ Tip

To use the `xs` CLI, you must first log on to an SAP HANA system as the operating-system user `<SID>adm` and then log on to the `xs` CLI, for example, as a user with `XSA_ADMIN` privileges. For more information about how to log on to the `xs` CLI, see *Related Information* below.

- The `xsa` command-line interface

A selection of utilities to help you maintain XS advanced run-time instances, for example, enable, disable, restart XS advanced, or maintain domain certificates, etc.

```
XSA {COMMAND} [--OPTIONS]
```

! Restriction

To use the `xsa` command, you must log on as the operating-system user `<SID>adm`.

Related Information

[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

[Maintaining XS Advanced Run-Time Instances with the XSA CLI \[page 1740\]](#)

[Maintaining the XS Advanced Run-time Environment with a Graphical User Interface \[page 1753\]](#)

12.2.2.1 Maintaining XS Advanced Run-Time Components with the XS CLI

Use the `xs` command-line interface (CLI) to manage the XS advanced run-time environment.

SAP HANA provides a command-line interface that enables you to maintain not only the applications that are deployed to the XS advanced run-time environment, but also specific elements of the run-time environment itself, for example, the components that enable it, and the users who access and use it. For example, you can use the XS CLI to maintain and manage the following components:

- Logon and setup
- XS advanced applications, routes, and tasks
- Organizations and spaces
- Domains and certificates
- Services
- Build packs, run-time environments, and the blob store
- XS advanced users
- Security aspects
- XS advanced application deployment, installation, and life cycle

→ Tip

You can install and use the `xs` command-line interface tools over a secure connection from a remote machine. For more information about where to find and how to download the `xs` command-line interface tools, see *The XS Command-Line Interface* in *Related Information*.

Related Information

[The XS Command-Line Interface \[page 1653\]](#)

[Logging on and Getting Started with the XS CLI \[page 1656\]](#)

12.2.2.1.1 The XS Command-Line Interface

A list of all the categories and areas covered by the `xs` command-line interface (CLI).

The `xs` command-line interface is a set of tools that enable the administration of the XS advanced run time from the command-line. The `xs` CLI is automatically installed on the SAP HANA system during installation of XS Advanced. However, you can also install the `xs` CLI tools on a remote machine and log on to the XS advanced run-time environment from the remote machine using a secure connection.

→ Tip

The XS advanced command-line client is available for download from SAP Service Marketplace for those people with the required S-User ID. Alternatively, you can find it on the SAP HANA installation media.

Downloading the `xs` CLI from the Service Marketplace

Users with the required S-User ID can download the `xs` CLI from the SAP Service Marketplace:

► [Service Marketplace](#) ► [Software Downloads \[Downloads\]](#) ► [SUPPORT PACKAGES & PATCHES](#) ► [By Alphabetical Index \(A-Z\)](#) ► [H](#) ► [SAP HANA PLATFORM EDITION](#) ⌵:

- [SAP HANA PLATFORM EDITION 1.0](#) ► [XS ADVANCED RUNTIME](#) ► [XS RUNTIME 1](#) ⌵
- [SAP HANA PLATFORM EDITION 2.0](#) ► [XS RUNTIME 1](#) ⌵

Downloading the `xs` CLI from the SAP HANA Media

The `xs` CLI is also available for download from the SAP HANA installation media. You can find the Zip archive `xs.onpremise.runtime.client*` in the following location on the SAP HANA media:

`DATA_UNITS/XSA_CLIENT_10/xs.onpremise.runtime.client_<platform>-<version>.zip`

→ Tip

Extract the contents of the Zip archive to the desired location. On Unix machines, the default location is the directory `/bin`.

Usage

```
xs <command> [<ARGUMENTS>] [<OPTIONS>]
```

To display information about a specific `xs` command:

```
xs help <command>
```

To display a list of all available `xs` commands, along with information about environment variables, and global options:

```
xs help -a
```

XS Command Overview

Command Category	Description
Logon and setup	User logon, view user-organization (and space) targets, set API URLs
Application Management	Maintain SAP HANA XS applications: list, deploy, start, stop, stage, [...]
Services Management	Maintain SAP HANA XS services: list, create, delete, bind, update, [...]

Command Category	Description
Organizations	Maintain user organizations: create, list, rename, delete, [...]
Spaces	Manage user spaces: create, list, rename, delete, [...]
Domains	Manage XS advanced domains: create, list, delete, set certificates, [...]
Certificates	Manage XS advanced certificates: set, unset, list trusted certificates
Routes	Maintain application routes: create, list, map, unmap, delete, [...]
Build Packs	Maintain application build-packs: create, list, update, rename, delete, [...]
Run-time Environments	Maintain XS run-times: create, list, display information, search, update, delete, [...]
Tasks	Maintain XS application-related tasks: list, run, cancel
User Administration	Maintain SAP HANA users: create, list, purge, delete, set/unset organizations, spaces, roles, [...]
Administration	Maintain XS application traces and backups
Tenant Databases	Manage and maintain tenant databases and their mapping to organizations and spaces
Configuration	Set and maintain environment variables and groups
Blob Store	Manage and maintain the contents of the blob store
Advanced	Retrieve and display OAuth tokens for the current session
Other Commands	Display information about the SAP HANA XS version, CLI, and system
Plug-ins	Additional commands as plug-ins; install/remove product components, deploy multi-target applications (MTA)

Related Information

[SAP Software Download Center \(Logon required\)](#) 

[Logging on and Getting Started with the XS CLI \[page 1656\]](#)

12.2.2.1.2 Logging on and Getting Started with the XS CLI

Get to know the basic commands available with the `xs` command-line interface.

In this section, you learn how to use the `xs` command-line interface (CLI) to perform the following tasks:

- Display a list of all `xs` commands
- Display details of a specific `xs` command
- Log on to (and out of) an instance of SAP HANA XS, advanced model

→ Tip

For information about predefined SAP HANA XS advanced users, for example, system and technical users, see *User Administration and Authentication in SAP HANA XS Advanced* in the *SAP HANA Security Guide*.

Getting Help

The `xs` command-line interface provides comprehensive help for each `xs` command. To display an overview of all available `xs` commands included with the `xs` CLI, use the help command as shown in the following example:

```
$ xs help -a
```

To display usage details of a specific `xs` command, add the name of the command to the help command as shown in the following example:

```
$ xs help <command>
```

To display details of the current version of the `xs` command-line client, use the `version` command as shown in the following example:

```
$ xs version
```

Logging on

During logon the `xs` CLI tries to connect to the so-called API URL of the XS Controller. To find out the current API URL, log on to the SAP HANA system as the operating-system user `<SID>adm` and run the following command:

```
$ xs-admin-login --api  
https://api.example.org:31030
```

You can use the URL displayed in the command output to log on to the XS advanced instance with the `xs CLI` from a remote session:

→ Tip

If you do not specify any user name in the `xs login` command, the `xs CLI` prompts you for the logon credentials you want to use.

```
$ xs login -a https://api.example.org:31030
USERNAME> XSA_ADMIN
PASSWORD> *****
Authenticating...
ORG: myorg
```

By default, the logon process targets the organization that was created during installation, for example, `myorg`. The first time you log on to XS advanced, you are asked to specify the organizational space you want to log on to from a list of available spaces, as illustrated in the following example:

```
Existing spaces:
0. PROD
1. SAP
SPACE> 0
SPACE: PROD
```

Choose the space you want to log on to by entering the corresponding number in the list, for example, `0` to select `PROD` as the target logon space. After successful logon, XS advanced displays a short summary of the logon details, as illustrated in the following example:

```
API endpoint: https://api.example.org:31030 (API version: 1)
User:        XSA_ADMIN
Org:         myorg
Space:       PROD
```

If used remotely, the `xs` client needs a trust certificate to establish a secure connection. For more information about obtaining this trust certificate, see [Trust Certificates \[page 1658\]](#).

Alternatively, if you fully trust the network connection to the server, you can skip the SSL validation process, as shown in the following example:

⚠ Caution

It is strongly recommended **not** to skip SSL validation in a production environment.

≡ Sample Code

```
xs login -a <API URL> -u <username> --skip-ssl-validation
```

→ Tip

If you already logged on to an SAP HANA system as the operating-system user `<SID>adm`, you can use the command `xs-admin-login` (without any parameters or options) as a shortcut to log on to the XS advanced run time as user `XSA_ADMIN`. The `xs-admin-login` command automatically sets up the `xs CLI` and logs you on as `XSA_ADMIN` after prompting for the corresponding password.

Trust Certificates

If you log on to XS advanced from a remote machine, the `xs` CLI requires a trusted certificate to establish a secure connection, unless `https` is disabled or a certificate signed by a well-known CA is used at the XS Controller. For your convenience, a public trust certificate is stored on the SAP HANA system at the following location:

```
/hana/shared/<SID>/xs/controller_data/controller/ssl-pub/default.root.crt.pem
```

The certificate file can be copied to the client, where it can then be consumed by using the `--cacert` option, as illustrated in the following example:

```
$ xs login -a https://api.example.org:8080 --cacert <PATH>
```

`<PATH>` indicates the location on the client machine of the file `default.root.crt.pem` which contains the certificate to use to establish the trusted connection.

→ Tip

Neither of these steps is necessary if a certificate signed by a well-known CA is available. For more information about certificates and trusted connections, see *Maintaining Domains* in *Related Information*.

Logging off

To log out of the `xs` command-line interface client session, use the `logout` command as shown in the following example:

```
$ xs logout
```

Related Information

[Maintaining Trust Certificates in XS Advanced \[page 1706\]](#)

[Maintaining Domains in XS Advanced \[page 1708\]](#)

[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

12.2.2.1.3 Displaying the System Overview

Display an overview of the current configuration and status of the main XS advanced components.

If you are logged on to the XS advanced platform with the `xs` command-line interface, you can display an overview of all the most important components of the XS advanced instance, for example, the Execution

Agents and the XS Advanced Controller, as illustrated in the following example of the `xs system-info` command:

Output Code

Displaying XS Advanced Information

```
$ xs system-info
Getting system infrastructure information...
```

! Restriction

To execute the `xs system-info` command, you need XS advanced administrator permissions.

The output generated by the `xs system-info` command includes information about the following XS advanced components:

- [Execution Agents \[page 1659\]](#)
- [XS Controller Server Version \[page 1660\]](#)
- [Registered Service URLs \[page 1662\]](#)
- [Applications \[page 1662\]](#)
- [Organizations And Spaces \[page 1663\]](#)

XS Advanced Execution Agents

The Execution Agents section of the information displayed by the `xs system-info` includes details of the execution agents that are currently connected to the XS advanced Controller, as displayed in the following example output:

Output Code

`xs system-info` command output

```
Execution agents:
-----
1. host1.example.org:53648
   created at Dec 14, 2017 9:10:23 AM
   port range 50000-50499
   os.arch amd64
   os.name Linux
   java.vendor SAP AG
   java.version 1.8.0_144
   os.version 3.12.62-60.62-default
   sun.arch.data.model 64
   version v1.0.72
2. host2.example.org:53574
   created at Dec 14, 2017 9:11:24 AM
   port range 50000-50499
   os.arch amd64
   os.name Linux
   java.vendor SAP AG
   java.version 1.8.0_144
   os.version 3.12.62-60.62-default
   sun.arch.data.model 64
   version v1.0.72
```

In this example, two execution agents on "host1.example.org" and "host2.example.org" are connected to schedule application instances on. For each execution agent, additional information is displayed like:

- The host on which the Execution Agent is running
- The date and time when the Execution Agent was started
- The port range the Execution Agent uses for applications on its host
- The platform and version of the SAP Java Virtual Machine (JVM) that is running the Execution Agent
- The version of the Execution Agent

→ Tip

If you need to scale out your application, you can add additional execution agents to an SAP HANA system by adding additional SAP HANA hosts with role "xs_worker".

XS Advanced Controller

The command `xs system-info` displays information about the XS Controller, as shown in the following example:

Output Code

`xs system-info` Command Output

```
Controller server version information:
-----
name                XS Controller
support             http://service.sap.com/message
build              v1.0.72
api version         1
software version    1.0.72.0
content version     1.0.72.0
state               READY
description         SAP XS Runtime on premise
controller endpoint https://example.org:30030
authorization endpoint https://example.org:30032/uaa-security
accept encoding     gzip, x-gzip
usage              apps: 8, routes: 8, services: 16
                   instances: 0 starting, 6 running, 3 stopped, ...
                   1 crashed, 0 timed out
database type       HANA_MULTI
database info       HDB 2.00.022.00.1511184640
runtime database    SYSTEMDB
```

The following table lists and explains the information that the `xs system-info` command displays about the XS Controller.

XS Controller Information

Information	Description
name	The name of the XS advanced component
support	The URL to use to contact a support team

Information	Description
<code>build</code>	The built version of the XS Controller
<code>api version</code>	The version of the REST interface provided by the XS Controller
<code>software version</code>	The installed version of the XS advanced system components.
	<p>i Note</p> <p>In a standard installation, <code>software version</code> and <code>content version</code> must match.</p>
<code>content version</code>	The installed version of the XS advanced core applications package (for example, <code>deploy-service</code> and <code>product-installer</code>).
	<p>i Note</p> <p>In a standard installation, <code>software version</code> and <code>content version</code> must match.</p>
<code>state</code>	<p>The current state of the XS advanced run-time, which can be one of the following:</p> <ul style="list-style-type: none"> • <code>STARTUP</code> XS advanced is currently starting • <code>READY</code> XS advanced is installed correctly • <code>INSTALLATION_PENDING</code> The initial XS advanced installation is not yet complete • <code>UPDATE_PENDING</code> The XS advanced update is not yet complete
<code>description</code>	A short summary of the type of installed component, for example, <code>SAP XS Runtime on premise</code>
<code>controller endpoint</code>	The REST end point for requests to the XS Controller API (also known as the "API URL")
<code>authorization endpoint</code>	The REST end point for authorization requests
<code>usage</code>	Information about the number of applications and application instances and their current status, for example: <code>starting</code> , <code>running</code> , <code>stopped</code> , etc.
<code>database type</code>	<p>The type of the underlying SAP HANA database installed; this can be either of the following:</p> <ul style="list-style-type: none"> • <code>HANA_MULTI</code> SAP HANA multiple database-container system • <code>HANA_SINGLE</code> SAP HANA single database-container system
<code>database info</code>	The version of the underlying SAP HANA database

Information	Description
runtime database	The database containing the persistence of the XS Advanced system components

XS Advanced Service URLs

The `xs system-info` command displays information about the services URLs registered by deployed XS advanced applications. XS advanced applications can maintain a service URL during installation to publish their application end point:

Output Code

`xs system-info` Command Output

Registered service URLs:

```
-----
deploy-service      https://example.org:51004
product-installer   https://example.org:51005
[...]
```

XS Advanced Applications Overview

The `xs system-info` command uses the "Applications" section to display statistics about the applications deployed in the XS advanced system and the current known status, as illustrated in the following example output:

Output Code

`xs system-info` Command Output

Applications:

```
-----
Deployed applications: 8
Running applications:  6
Stopped applications:  2
```

Organizations And Spaces in XS Advanced

The `Organizations` and `Spaces` section of the information displayed by the `xs system-info` command displays statistics about the distribution of applications across the configured organizations and spaces, as illustrated in the following example:

```
☰ Output Code
xs system-info Command Output

Organizations and spaces:
  Organization      Space      Number of apps      User:
-----
  myorg
  |-                PROD       0                    abcxs
  |-                SAP        8                    sapabcxs
```

In the example displayed, the organization "myorg" contains two spaces named "PROD" and "SAP". In addition to the number of applications deployed in each space, the information displayed also indicates the name of the operating-system user (for example, `<SID>xs`) used to start the XS advanced application.

Related Information

[Maintaining Services in XS Advanced \[page 1692\]](#)

[Maintaining the XS Advanced Run-time Environment with a Command-Line Interface \[page 1652\]](#)

12.2.2.1.4 Maintaining Organizations and Spaces in XS Advanced

Use the `xs` command-line interface to configure and maintain organizations and spaces in XS advanced.

The examples in this section show how to use the `xs` command-line interface (CLI) to perform the following tasks:

- [Navigating Through Organizations and Spaces \[page 1664\]](#)
- [Maintaining Spaces \[page 1666\]](#)
- [Maintaining Organizations \[page 1666\]](#)
- [Isolating Applications at the Operating-System Level \[page 1667\]](#)
- [Custom Operating-System Users \[page 1668\]](#)

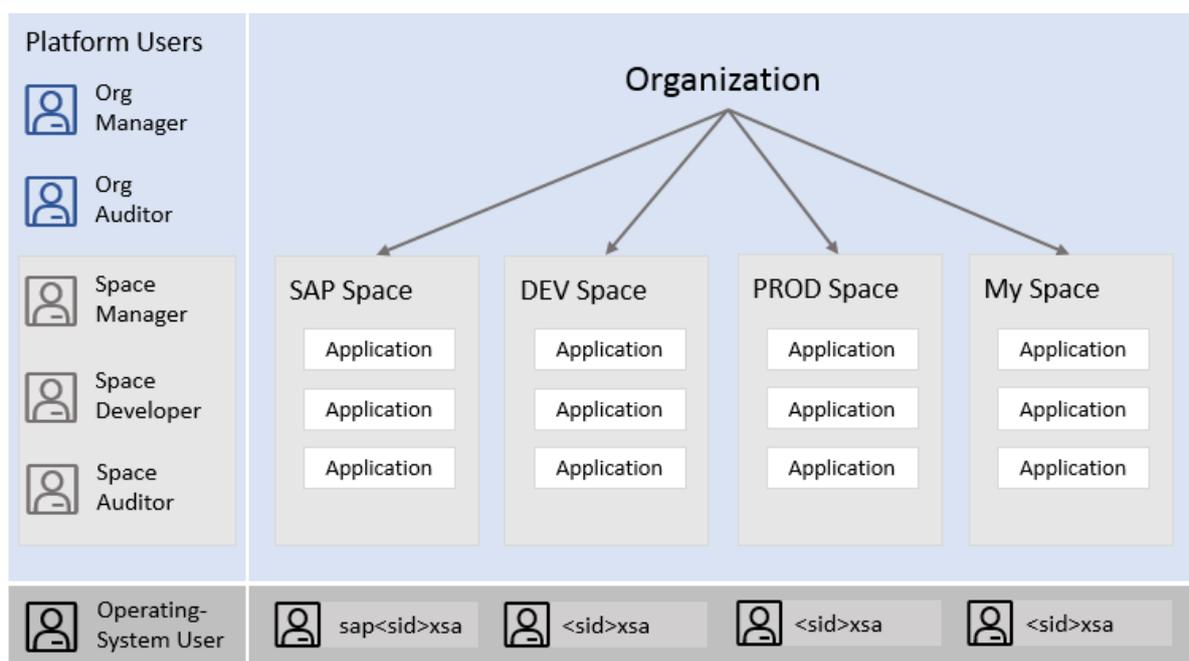
Overview

XS advanced applications can be grouped together by means of a hierarchy of so-called organizations and spaces. An organization forms a logical group of spaces. Information about application domains and server

certificates is also tied to organizations. In XS advanced, a “space” is not only a way to form a logical group of applications, it is also a trust zone. Applications cannot use any resources that were created in another space (for example, service instances), and application processes are also isolated at the space level by using operating-system users to start the application processes, as described in [Custom Operating-System Users \[page 1668\]](#) below.

→ Tip

See *Maintaining Platform Users in Related Information* for more information about user roles in organizations and spaces. For more information about the underlying security concepts for organizations and spaces, see the *SAP HANA Security Guide in Related Information* below.



Navigating Through Organizations and Spaces in XS Advanced

During the installation of XS advanced, one organization and two spaces are created by default:

- An organization with a custom name
- A space named `SAP` which contains the default XS advanced applications provided by SAP
- A space for custom development or deployment; by default, this space is named `PROD`.

To display a list of all organizations configured in an XS advanced run-time instance, run the following command, which in this example, shows the organizations `acme` and `corp`:

```
$ xs orgs
Getting orgs...
name
-----
acme
```

```
corp
```

To view the spaces configured in the currently targeted organization, use the command `xs spaces`, as shown in the following example:

```
$ xs spaces

Getting spaces in org acme as XSA_ADMIN ...

name      user
-----
DEV       xsaxsa
PROD      xsaxsa
SAP       sapxsaxsa
```

→ Tip

The space `name` is shown along with the name of the operating-system `user` whose account is used to start applications in the specified space, for example, the user `xsaxsa` starts applications in spaces `DEV` and `PROD`, and the user `sapxsaxsa` starts applications in space `SAP`.

You can set the target for the `xs` CLI to a specific organization and space by using the `xs target` command, as shown in the following example:

```
$ xs target -o <org-name> -s <space-name>
```

When you have set the target organization and space, you can use the `xs apps` command to display a list of the applications currently running there:

```
$ xs apps

Getting apps in org "acme" / space "SAP" as XSA_ADMIN...
Found apps:

name                    state    instances memory    disk          urls
-----
auditlog-db             STOPPED  0/1      16.0 MB  <unlimited>  <none>
auditlog-server         STARTED  1/1      256 MB   <unlimited>  https://acme.org:
auditlog-broker         STARTED  1/1      64.0 MB  <unlimited>  https://acme.org:
deploy-service          STARTED  1/1      280 MB   <unlimited>  https://acme.org:
component-registry-db  STOPPED  0/1      16.0 MB  <unlimited>  <none>
product-installer       STARTED  1/1      256 MB   <unlimited>  https://acme.org:
auditlog-odata          STARTED  1/1      128 MB   <unlimited>  https://acme.org:
auditlog-ui             STARTED  1/1      64.0 MB  <unlimited>  https://acme.org:
xsa-admin-backend       STARTED  1/1      128 MB   <unlimited>  https://acme.org:
xsa-admin               STARTED  1/1      128 MB   <unlimited>  https://acme.org:
```

To display details of the currently targeted organization and space, use the command `xs target` without any parameters or options, as illustrated in the following example:

```
$ xs target

API endpoint: https://xsa.acme.com:30030 (API version: 1)
User:        XSA_ADMIN
Org:         acme
Space:       SAP
```

To change to a specific target organization and space, use the `xs target` command with the options `-o` `<ORG>` `-s` `<SPACE>`, as illustrated in the following example:

```
$ xs target -o acme -s DEV
```

Maintaining Spaces in XS Advanced

In some advanced use cases, the default space setup is not sufficient to meet the needs of your environment. For example, if you want to separate applications currently in development from applications that have been tested (so that they run in different spaces), you can create additional spaces, for example, named `DEV` and `TEST`. Similarly, you can separate different development teams by ensuring that they work in different spaces. To create a new space, use the command `xs create-space`, as shown in the following example:

```
$ xs create-space DEV
Creating space DEV in org acme as XSA_ADMIN...
OK
```

The command shown in the example above will create an empty space named `DEV` within the currently targeted organization `acme`.

→ Tip

The `xs create-space` command does not automatically switch target to the new space `DEV`; you must manually set the new space as the new target, for example, with the command `xs target -s <SPACE>`.

To delete an existing space, use the command `xs delete-space`, as shown in the following example:

```
$ xs delete-space DEV
Really delete the space "DEV" in org "acme"? (y/n) > y
Deleted space "DEV".
```

You can change the name of an existing space, too, as illustrated in the following example, which shows how to change the name of the space `DEV` to `TEST`:

```
$ xs rename-space DEV TEST
Renaming space "DEV" in org "acme"...
OK
```

Maintaining Organizations in XS Advanced

Organizations are particularly useful if you want to expose one group of applications by means of a different domain name to the one used by another group of applications. Hence, organizations can be used to reflect the organizational structure of your company or corporation. In order to create a new, empty organization, use the command `xs create-org`:

```
$ xs create-org myorg
```

OK

→ Tip

The `xs create-org` command does not automatically switch target to the new organization `myorg`; you must manually set the new organization as the new target, for example, with the command `xs target -o <ORGANIZATION>`.

To delete an existing organization, use the command `xs delete-org`, as shown in the following example:

```
$ xs delete-org myorg
Deleting org "myorg"...
Really delete the org "myorg"? (y/n) > y
Deleted org "myorg".
```

⚠ Caution

The `xs delete-org` command also deletes **all spaces** contained in the deleted organization along with all related resources such as applications, service instances, and routes. Note that all applications running in any of the deleted organization's spaces must be stopped before you delete the organization.

You can change the name of an existing organization, too, as illustrated in the following example, which shows how to change the name of the organization "acme" to "test":

```
$ xs rename-org acme test
OK
```

Isolating XS Advanced Applications at the Operating-System Level

Instances of running application can be isolated at the space level. This means that processes of application instances started in the same space are run as the same operating system user. You can also specify that application processes from different spaces are executed as different operating system users. In this way, application processes from an isolated space are not allowed to access the process environment and the file system of other applications or the SAP HANA system itself.

During installation, two restricted operating system users are automatically created by `hdb1cm` and assigned to default XS advanced spaces. These users are automatically maintained by `hdb1cm` if an additional `xs_worker` host is added to the SAP HANA system or during any rename operation. The following table shows the mapping of the default operating system users to XS advanced spaces;

Default Username-Space Mapping in XS Advanced

User Name	Space Name
<code>sap<SID>xsa</code>	SAP
<code><SID>xsa</code>	Custom space (for example, <code>PROD</code>); this is the default user name for newly created spaces

In this example, application processes in the space `SAP` run as user `sap<SID>xsa` while application processes in the space `PROD` will run as user `<SID>xsa`. User names are combined with the System ID (SID) of the SAP

HANA system. This is to prevent any clashes between user names when several SAP HANA systems are installed on the same host. To change the operating system user mapping for a particular space, use the command `xs update-space`, as shown in the following example:

```
$ xs update-space <space-name> -u <OS-user-name>
```

i Note

You cannot change the mapping between operating-system user and a space if applications are still running inside the targeted space. You must stop all applications that are running in a space before changing the operating-system user mapping.

To revert to the default settings for the mapping between space and operating-system user, use the `--unset-user` option as shown in the following example:

```
$ xs update-space <space-name> --unset-user
```

Custom Operating-System Users

The command `xs update-space` does not create the underlying operating system user. If a custom operating system user is required in addition to the default operating system users, the new custom operating-system user must be created first and on all SAP HANA hosts performing the role of `xs_worker` or `xs_standby`.

i Note

The operating-system user associated with a space only requires a restricted set of privileges.

The operating-system user associated with a space only requires a restricted set of privileges:

- The operating system user does **not** need any elevated or superuser privileges
- The operating system user must **not** be in the group `sapsys`
- No secondary groups are required
- No `home/` directory is required
- No user logon is required

In addition, if you use a SAP HANA version older than HANA 2.0 SPS 02 Rev. 1, it is necessary to adapt the `sudo` configuration file `/etc/sudoers`. For more details about how to set up a new operating system user and adapt the `sudoers` configuration files, see also SAP Note [2243156](#).

Related Information

[Maintaining Platform Users in XS Advanced \[page 1671\]](#)

[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

[The XS Command-Line Interface \[page 1653\]](#)

12.2.2.1.5 Maintaining Organization Quotas

Set limits for application instances for XS advanced organizations and maintain organizational quota plans.

Organization quotas are named sets of restrictions applied to an organization by an organization manager. For example a quota might limit an organization to create up to 10 application instances. Because organization quotas are a collection of such individual limits they are also referred to as quota plan, quota definition or simply quotas. The limits supported by organization quotas are shown in the table below.

XS Advanced Organization Limits

Limit	Description
App instance limit	The total number of instances of started applications allowed in the specified organization. Instances of stopped applications do not count toward this instance limit. i Note "-1" represents an unlimited amount. (Default: "-1")
Max instances per application	The total number of running instances allowed per application in the specified organization. i Note "-1" represents an unlimited amount. (Default: "100")
Service instances limit	The total number of service instances allowed in the specified organization. i Note "-1" represents an unlimited amount. (Default: "-1")

When a new organization is created, an organization quota named "default" is automatically assigned to it; this is referred to as the "default quota". Although the default quota cannot be deleted, the values can be modified by an administrative user. Only an XS advanced controller administrator can use the quota commands to create, update, or delete an organization quota. An organization manager can use the quota commands to assign a quota to an organization.

Creating a Quota Plan for an Organization

The following example shows how to use the `xs create-quota` command to set a quota plan for the current target organization:

```
xs create-quota quota1 -a 10 -ma 5 -s 4
```

The command illustrated in the example above creates a quota named "quota1" that limits the total amount of started application instances (-a) to "10"; the maximum amount of running instances per application (-ma) to "5"; and the total amount of service instances (-s) to "4". Where no option is specified, the default value for the respective limit is assigned.

Reading an Organization Quota Plan

The following example shows how to use the `xs quotas` command to display a list of all existing organization quota plans:

xs quotas

Getting Quota Definitions...

name	app instances	max instances per app	service instances
default	unlimited	100	unlimited
quotat1	10	5	4

To display details of a specific quota, use the `xs quota <QUOTA_NAME>` command, as illustrated in the following example:

xs quota quotat1

Property	Value
Name	quotat1
Service instances limit	unlimited
App instance limit	unlimited
Max instances per app	100

Assigning a Quota Plan to an Organization

The following example shows how to use the `xs set-quota` command to assign a quota plan to a specific organization:

```
xs set-quota org1 quotat1
```

The command illustrated in the example above assigns the quota named "quotat1" to the organization named "org1". Assigning of a quota to an organization will fail if the limits set in the assigned quota are already being broken by the organization. For example if the organization has 100 application instances, then a quota with a total application instance limit of "10" cannot be assigned.

i Note

A quota can be assigned to multiple organizations. Any changes made to the quota will affect all the organizations it is assigned to.

To view the quota definition assigned to an organization use the command "xs org". The example below demonstrates an elaborate use:

xs org org1

Getting info for org "org1" as user "XSA_ADMIN"...

```
name:      org1
domains:   example.org
quota:     quotat1
spaces:    spacel
```

```
created: May 30, 2018 8:55:36 AM
updated: May 30, 2018 8:55:36 AM
```

Deleting a Quota Plan

The following example shows how to use the `xs delete-quota` command to unassign and remove a quota plan:

```
xs delete-quota quota1
```

The command illustrated in the example above deletes the quota named "quota1".

! Restriction

An organization quota can only be deleted if it is not assigned to any organization. The quota named "default" cannot be deleted.

Related Information

[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

[Maintaining Organizations and Spaces in XS Advanced \[page 1663\]](#)

12.2.2.1.6 Maintaining Platform Users in XS Advanced

Use the `xs` command-line interface to maintain XS advanced users.

Users who have access to the XS advanced platform (so-called platform users) can be sorted into the following categories:

- Platform administrators
For example, `XSA_ADMIN`, who are allowed to perform any platform operation in any organization and space.
- Restricted users
For example, developers and organization or space manager, who have restricted permissions in specific organizations or spaces. A role model exists for restricted platform users, which is defined in the *SAP HANA Security Guide*. For more details about this user model, see *Related Information*.

New XS advanced platform users are created by assigning predefined platform role collections to existing SAP HANA users, for example, with the `xs` CLI, as described in this section. The following table shows the XS advanced role collections that exist for platform users:

XS Advanced Role Collections

Role Collection	Description
<code>XS_CONTROLLER_ADMIN</code>	The platform administrator
<code>XS_CONTROLLER_USER</code>	Enables a restricted user who has refined permissions in organizations and spaces
<code>XS_CONTROLLER_AUDITOR</code>	Enables a restricted read-only user who has refined permissions on organizations and spaces
<code>XS_USER_ADMIN</code>	Enables a user who can create and maintain other users
<code>XS_AUTHORIZATION_ADMIN</code>	Enables a user who can view, create and maintain role collections

Viewing Assigned Role Collections

To view the role collections assigned to a specific user, use the command `xs assigned-role-collections`, as illustrated in the following example:

```

🔍 Output Code

Role Collections Assigned to a User

$ xs assigned-role-collections XSA_ADMIN

role collection      description
-----
XS_CONTROLLER_ADMIN  Authorizations for XS controller admin
XS_AUTHORIZATION_ADMIN  Authorizations for XS role builder
XS_USER_ADMIN        Admin authorizations for XS user management
AUDITLOG_VIEWER
  
```

In the example output above, the default admin user `XSA_ADMIN` has a role collection `XS_CONTROLLER_ADMIN` that enables the administrator to perform **any** task in **any** XS advanced organization and space. In addition, `XSA_ADMIN` has the role collections `XS_AUTHORIZATION_ADMIN` and `XS_USER_ADMIN`, which enable the administrator user to create other users and maintain and assign role collections. `XSA_ADMIN` has an additional application role collection that provides the permissions required to view audit logs.

Assigning Role Collections

The following sections describe how to create new administrators and restricted users by assigning the corresponding role collection and setting permissions that apply for particular organizations and spaces.

Creating XS Advanced Administrator Users

To create a new XS advanced administrator user, an existing SAP HANA user in the system database is required. The following example shows how to use the `hdbsql` utility to create an SAP HANA user. You need to log on to the SAP HANA system with `hdbsql` as a user with the `User.Admin` role and execute the following SQL command:

```
=> CREATE USER XSA_ADMIN2 PASSWORD "Welcome1"
0 rows affected (overall time 25.359 msec; server time 23.375 msec)
```

If you are logged to the XS advanced platform as an XS advanced administrator (for example, `XSA_ADMIN`), you can grant the new XS advanced administrator user (`XSA_ADMIN2`) the necessary controller roles `XS_CONTROLLER_ADMIN`, `XS_USER_ADMIN`, and `XS_AUTHORIZATION_ADMIN`, for example, using the command `xs assign-role-collection`:

Sample Code

```
$ xs assign-role-collection XS_CONTROLLER_ADMIN XSA_ADMIN2
Assigning role collection "XS_CONTROLLER_ADMIN" to user "XSA_ADMIN2"...
OK

$ xs assign-role-collection XS_USER_ADMIN XSA_ADMIN2
Assigning role collection "XS_USER_ADMIN" to user "XSA_ADMIN2"...
OK

$ xs assign-role-collection XS_AUTHORIZATION_ADMIN XSA_ADMIN2
Assigning role collection "XS_AUTHORIZATION_ADMIN" to user "XSA_ADMIN2"...
OK
```

The `XSA_ADMIN2` can be now used to perform administrator tasks on the XS Advanced platform.

! Restriction

To log on to the new XS advanced administrator account (`XSA_ADMIN2`), you must first change the initial password. For example using tools included in SAP HANA Cockpit or SAP HANA Studio.

The `xs` CLI points you to the XS advanced authorization endpoint that you can access with a Web Browser, where you can change the initial password for the new XS advanced platform administrator user. Alternatively, you can use the command `hdbsql` to set the new password.

Creating XS Advanced Restricted Users

To create a new XS advanced **restricted** user, an existing SAP HANA user in the system database is required. The following example shows how to use the `hdbsql` utility to create an SAP HANA user. You need to log on to the SAP HANA system with `hdbsql` as a user with the `User.Admin` role and execute the following SQL command:

```
=> CREATE USER NEW_XSA_USER PASSWORD "Welcome1"
0 rows affected (overall time 26.582 msec; server time 24.367 msec)
```

If you are logged to the XS advanced platform on as an XS advanced administrator (for example, `XSA_ADMIN`), you can grant the new XS advanced user (`NEW_XSA_USER`) the necessary controller roles `XS_CONTROLLER_ADMIN`, for example, using the command `xs assign-role-collection`:

```
$ xs assign-role-collection XS_CONTROLLER_USER NEW_XSA_USER
Assigning role collection "XS_CONTROLLER_USER" to user "NEW_XSA_USER"...
OK
```

With the assignment of the role collection `XS_CONTROLLER_USER`, the new user can log on to the XS advanced platform but still does not have any privileges. XS Controller roles need to be assigned to the restricted XS advanced user at the organization and space level, which is described in the following sections.

Setting Roles and Permissions at the Organization Level

The `xs` command-line client includes the command `xs set-org-role`, which you can use to grant organization-specific roles to users. For example, to grant the user `XSA_USER` the `OrgManager` role in the organization “acme”, run the following command:

```
$ xs set-org-role XSA_USER acme OrgManager
OrgManager Adding role 'OrgManager' to user NEW_XSA_USER in org "acme" ...
OK
```

To set user permissions at the organization level, you must log on to XS advanced with either administrator privileges or as a restricted user with the role `OrgManager` in the organization where you want to grant user roles and privileges. The following table lists the roles you can grant to an organization user:

XS Advanced Organization Roles

Organization Role	Privileges
<code>OrgManager</code>	Create and modify spaces in an organization Assign roles and privileges to other organization users
<code>OrgAuditor</code>	Browse through spaces inside an organization

To remove a role already assigned to an organization user, use the `unset-org-role` command, as illustrated in the following example:

```
$ xs unset-org-role XSA_USER acme OrgManager
```

Setting Roles and Permissions at the Space Level

The `xs` command-line client includes the command `xs set-space-role`, which you can use to grant space-specific roles to users. For example, to grant the user `NEW_XSA_USER` the `SpaceManager` role for a space called “DEV” in an organization named “acme”, run the following command:

```
$ xs set-space-role XSA_USER acme DEV SpaceManager
Adding role 'SpaceManager' to user XSA_USER in space "DEV" in org "acme" ...
OK
```

To set user permissions at the space level, you must log on to XS advanced with either administrator privileges or as a restricted user with the role `SpaceManager` in the space where you want to grant user roles and privileges. The following table lists the roles you can grant to a space user:

XS Advanced Organization Roles

Space Role	Privileges
<code>SpaceManager</code>	Modify a space Assign roles and privileges to other space users
<code>SpaceDeveloper</code>	Modify a space
<code>SpaceAuditor</code>	Browse through any space where you have the <code>SpaceAuditor</code> role

To remove a role already assigned to a space user, use the `unset-space-role` command, as illustrated in the following example:

```
$ xs unset-space-role XSA_USER acme DEV SpaceManager
```

Viewing Available Role Collections

You can view the available XS advanced role collections by using the command `xs role-collections`. The command output shows details of role collections for the XS advanced platform as well as the role collections created for XS advanced applications, as illustrated in the following example:

Sample Code

```
$ xs role-collections
```

```
Getting role collections as user "XSA_ADMIN"...
```

```
role collection      description
-----
AUDITLOG_VIEWER
XS_AUTHORIZATION_ADMIN  Authorizations for XS role builder
XS_AUTHORIZATION_DISPLAY Authorizations for XS role viewer
XS_USER_ADMIN           Admin authorizations for XS user management
XS_USER_DISPLAY         Display authorizations for XS user management
XS_USER_PUBLIC          Default authorizations for XS user
XS_MONITOR_ADMIN        Authorizations for XS monitoring management
XS_MONITOR_DISPLAY      Authorizations for XS monitoring display
XS_SUBSCRIPTION_ADMIN   Authorizations for XS subscriptions management
XS_SUBSCRIPTION_DISPLAY Authorizations for XS subscriptions display
XS_TENANT_ADMIN         Authorizations for XS tenants management
XS_TENANT_DISPLAY       Authorizations for XS tenants display
XS_CONTROLLER_ADMIN     Authorizations for XS controller admin
XS_CONTROLLER_USER      Authorizations for XS controller user
XS_CONTROLLER_AUDITOR   Authorizations for XS controller auditor
```

Related Information

[Building Roles for XS Advanced Applications \[page 1725\]](#)

[Maintaining Organizations and Spaces in XS Advanced \[page 1663\]](#)

[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

12.2.2.17 Displaying Application Information in XS Advanced

Use the `xs` command-line interface to deploy and maintain XS advanced applications.

After navigating to a particular organization and space, you can explore the resources contained in a space. A space contains different types of resources, for example: applications, service instances, and routes. The information in the following sections focuses on application resources and aims to help you perform the following application-related tasks in an XS advanced space:

- [Displaying Deployed XS Advanced Applications in a Space \[page 1676\]](#)
- [Displaying Details of an Individual XS Advanced Application \[page 1678\]](#)
- [Displaying XS Advanced Application Logs \[page 1680\]](#)
- [Displaying XS Advanced Application Events \[page 1682\]](#)
- [Displaying XS Advanced Application Files \[page 1682\]](#)

Displaying Deployed XS Advanced Applications in a Space

When you set a space as a particular “target”, you can use the `xs` CLI to display information about the XS advanced applications running in the target space, as shown in the following example:

Output Code

```
$ xs apps

Getting apps in org "acme" / space "SAP" as XSA_ADMIN...
Found apps:
name                requested state instances memory  disk          alerts  urls
-----
auditlog-db         STOPPED  0/1      16.0 MB <unlimited>   <none>
auditlog-server     STARTED  1/1      256 MB  <unlimited>   https:
auditlog-broker     STARTED  1/1      64.0 MB <unlimited>   https:
deploy-service      STARTED  1/1      280 MB  <unlimited>   https:
component-registry-db STOPPED  0/1      16.0 MB <unlimited>   <none>
product-installer   STARTED  1/1      256 MB  <unlimited>   https:
auditlog-odata      STARTED  1/1      128 MB  <unlimited>   https:
auditlog-ui         STARTED  1/1      64.0 MB <unlimited>   https:
xsa-admin-backend   STARTED  1/1      128 MB  <unlimited>   https:
xsa-admin           STARTED  1/1      128 MB  <unlimited>   https:
demo-app            STARTED  2/3      256 MB  <unlimited>   1/1/1 https:
demo-app2          STARTED  0/1      256 MB  <unlimited>   DOWN! https:
```

The information displayed in the example above includes the following details:

XS Advanced Application Details

Property	Description
name	The names of the XS advanced applications deployed in the current organization/space
requested state	<p>The targeted state of the application, for example, <code>STARTED</code> or <code>STOPPED</code>. This state results from a request by an administrator to start or stop an application.</p> <p>→ Tip</p> <p>The <code>requested state</code> does not indicate the actual state of any application instances. The actual status of an application is displayed in the <code>instances</code> property.</p>
instances	<p>The actual status of an application, described by two numbers separated by a slash, for example, <code>0/1</code>:</p> <ul style="list-style-type: none">• The number of application instances that are currently running• The targeted number of running instances <p>→ Tip</p> <p>An application is fully started if the number of running instances matches the targeted number of running instances, for example, <code>1/1</code>.</p>
memory	The amount of memory assigned to the application
disk	The disk quota allocated for use by the application
alerts	Information about problems with the application. If there were crashes of application instances, three values are displayed, for example, " <code>1/1/1</code> ", which represent the number of crashes within different periods of time (short term: 5 minutes, mid-term: 1 hour, long-term: 1 day). If XS advanced cannot start an application, ' <code>DOWN!</code> ' is displayed.
urls	The Universal Resource Location (URL), defined in a "route" where the application is accessible

When reading the information displayed by the `xs apps` command, bear in mind the following details:

- For most applications, the expected number of application instances are running.
- Some applications are `STOPPED`, for example, at the request of an XS advanced administrator. Any application that performs a single task stops automatically as soon as it completes the assigned task. For example, the applications "`*-db`" in the example above (`auditlog-db` and `component-registry-db`) perform a one-off task such as initial database deployment.
- The application `demo-app2` is down and no further application instances will be started by XS advanced.
- Only 2 of the possible 3 instances of the application "`demo-app`" are started. There was a crash of an instance within the last 5 minutes. This could indicate a problem that needs further investigation.

Displaying Details of an XS Advanced Application Instance

To display detailed information about a specific instance of an XS advanced application, use the `xs app` `<appName>` command, as shown in the following example:

Output Code

```
$ xs app demo-app

Showing status and information about "demo-app"
name: demo-app
requested state: STARTED
instances: 1
memory: 128 MB
disk: <unlimited>
buildpack: sap_nodejs_buildpack
urls: https://acme.org:63055
created: Oct 17, 2017 2:46:40 PM
updated: Nov 28, 2017 11:14:04 AM
weighted instance uptime: 99.6 %
ALERTS: There were crashes!
  short term (last 5 m) 1
  mid term (last 1 h) 1
  long term (last 1 d) 1
  total 2

Instances of droplet 1 created at Oct 17, 2017 2:46:45 PM
index created state host port os user
-----
0 Oct 19, 2017 1:34:08 PM CRASHED acme.org 50012 xsaxsa

Instances of droplet 2 created at Nov 28, 2017 11:14:10 AM
index created state host port os user
-----
0 Nov 28, 2017 10:13:44 AM STOPPED acme01 50030 xsaxsa
1 Nov 28, 2017 11:14:48 AM RUNNING acme01 50030 xsaxsa
2 Nov 28, 2017 11:14:48 AM RUNNING acme01 50031 xsaxsa
3 Nov 28, 2017 11:14:48 AM CRASHED acme01 50033 xsaxsa
4 Nov 28, 2017 11:15:53 AM STARTING acme01 50032 xsaxsa
```

If alerts for an application are present, the `xs app` command shows detailed information about those alerts. For example, it shows the number of application crashes within short-, mid-, and long-term periods as well as the total number of crashes registered for an application.

For each application instance, the `xs app` command displays the following information:

- `index`
The index of the application instance used to reference it with further commands
- `created`
The time at which the application instance is created
- `state`
The current state of the application instance:
 - `STARTING`
The application instance has started, for example, at the request of an administrator, but is not yet accessible.
 - `RUNNING`
The application instance started successfully and is accessible.
 - `STOPPED`

The application instance stopped, for example, at the request of an administrator, and is not accessible.

- `CRASHED`

The application instance terminated unexpectedly.

→ Tip

You can use the command `xs delete-app-instances` to clean up instances of an application in the state `CRASHED`.

- `host`

The host on which the application instance is started

- `internal port`

The internal port of the application instance

→ Tip

The internal port is not the port on which the application is reachable by a business user; it is the number of the **internal** system port to which the XS advanced Platform Router forwards requests. The external end point for an application instance is exposed by the Platform Router.

- `os user`

The operating system user in whose account the processes of the specified application instances are started.

The output returned by the `xs app` command also includes information about the selected application's "technical" version and "droplet" version. Each time you stage an application as part of the deployment push to the XS advanced run-time platform, a so-called droplet is created which can be referenced by an index. Although it is often the case that only application instances of the most recent droplet are running, it is also important to understand that information about a certain number of previous droplets and application instances is kept for supportability reasons. In the example output above, two droplets of the same application are running, and in the command output, the droplets are indicated by "droplet 1" and "droplet 2", where `droplet 2` is the most recent version of the application.

- `droplet 2`

The most recent version of the application, as indicated by the creation time stamp, and the higher the droplet index the more recent the version. You can also see that two application instances are currently running and one application instance is still starting, which explains why the expected number of application instances has not yet been reached. One application instance crashed recently which caused the alert. One application instance of droplet 2 is `STOPPED`, which means that this instance of the application was started but the application processes are no longer running.

i Note

The contents of the `STOPPED` application instance's file system sandbox are retained so that support teams can analyze log and applications files at a later point in time, if necessary.

- `droplet 0`

The previous version of the application. The application instance with index 0 of droplet 1 is in state `CRASHED`, which means that the application was started before but terminated unexpectedly. Similarly to stopped applications, the file-system sandbox of crashed application instances is kept for the most recent instances so that support teams can analyze log and application files to find out the cause of the error.

→ Tip

Only the five most recently stopped or crashed instances are kept; any older application instances are deleted automatically.

Displaying XS Advanced Application Logs

To query the logs written by an XS advanced application (for example, output written to `stdout` and `stderr` as well as any access logs), run the `xs logs` command, as illustrated in the following example:

```
$ xs logs myAPP
```

When called without parameters, the `xs logs` command switches to "tailing" mode, where only newly written application logs are displayed in the system output, and the command prompt is returned when the user presses the keyboard combination: `CTRL` + `C`. To display the most recent entries of an application's logs, use the option `--recent` as shown in the following example:

```
$ xs logs myAPP --recent
```

To display all logs written by an application, use the option `--all` as shown in the following example:

```
$ xs logs myAPP -all
```

By default, the `xs logs` command displays the logs of **all** application instances. If you are only interested in the logs of a particular instance of an application, you can reference the application instance by including the command options `--droplet` and `--instance`. For example, if you want to display the logs of the crashed instance of the `demo-app` application in the example above, run the following command:

≡ Output Code

```
$ xs logs demo-app --all --droplet 1 --instance 0
11/27/17 4:04:57.122 PM [API] OUT Created app 'demo-app' [Org 'myorg', ...
11/27/17 4:04:57.360 PM [API] OUT Updated files for app 'demo-app' [Org...
11/27/17 4:04:57.390 PM [API] OUT Created droplet with id 1 of app 'demo-...
11/27/17 4:04:57.398 PM [API] OUT Staging Droplet with id 1 of app demo-a...
11/27/17 4:04:59.875 PM [STG/1] OUT Node.js
11/27/17 4:05:00.555 PM [STG/1] OUT Node.js buildpack version 3.3.2 [...]
11/27/17 4:05:02.066 PM [STG/1] OUT Copying SSL CA certificates...
...
```

→ Tip

Each log entry includes a time stamp, followed by the log source in brackets, followed by the log type and the actual log message. For more information about log sources and log types as well as how to filter for a specific type or source, see the following sections.

XS Advanced Application Log Sources

The following tables lists the sources used by XS advanced applications when writing log files:

XS Advanced Application Logfile sources

Log Source	Description
[API]	A log entry created by the XS advanced platform regarding this application
[STG/<Droplet index>]	A log entry created by the staging process regarding this application. The entry displays the droplet index as a suffix. For example, "STG/1" means that this staging process produced the droplet with index 1
[APP/<Droplet index>-<Instance index>]	The output of an application instance, either on <code>stdout</code> (prefixed with 'OUT') or on <code>stderr</code> (prefixed with 'ERR'). In the log type tag, the droplet and application instance is encoded in this case. In the example output above, "APP/1-0" means that the application instance with index 0 of droplet 1 dumped a log line to <code>stdout</code> .
[RTR]	The access log for this application created by the Platform Router (RTR)

It is possible to filter for certain log sources by specifying a log source (or multiple sources in a comma-separated list) in the option "`--source`" with the `xs logs` command, as shown in the following example:

Output Code

```
$ xs logs demo-app --all --source RTR
11/28/17 6:28:58.000 PM [RTR] OUT ##.##.209.139 - to acme.corp:63018 "GET...
11/28/17 6:28:58.000 PM [RTR] OUT ##.##.209.139 - to acme.corp:63018 "GET...
```

XS Advanced Application Log Types

A log source may contain log lines of different **types**. The following table lists the available log types:

XS Advanced Application Log File Types

Log Type	Description
OUT	A log entry was written to Standard Out (<code>stdout</code>)
ERR	A log entry was written to Standard Error (<code>stderr</code>)
ACC	The log line was written to the application instance's access log file
LOG	The log line was written in list log format to <code>stdout</code> and the logger starts with "/ Application"
SYS	The log line was written in list-log format to <code>stdout</code>

It is possible to filter the output to display only content from particular log **types**, for example, by specifying a log type (or multiple types in a comma-separated list) in the option "`--type`" with the `xs logs` command, as shown in the following example:

```
$ xs logs demo-app --all --type ERR
```

```
11/27/17 6:29:58.123 PM [APP/1-0] ERR Segmentation fault
```

You can combine options and log files as shown in the following example:

```
xs logs demo-app --recent --source APP --type ERR
11/27/17 6:29:58.123 PM [APP/1-0] ERR Segmentation fault
```

Displaying XS Advanced Application Events

The XS Controller saves important events that occur during the life time of an application. Viewing the events provide an overview of the changes made to an application from the point of view of the XS advanced platform:

Output Code

```
$ xs events deploy-service
```

```
Showing events for app "deploy-service" and user XSA_ADMIN
```

time	level	component	message

Dec 1, 2017			
11:26:00 AM	INFO	APPLICATION	Created app 'deploy-service' [Org ...
11:26:00 AM	INFO	SERVICE	Created service binding between app...
[...]			
11:26:04 AM	INFO	APPLICATION	Updated files for app 'deploy-service'...
11:26:04 AM	INFO	DROPLET	Created droplet with id 1 of app 'deploy-..
11:26:04 AM	INFO	DROPLET	Staging Droplet with id 1 of app deploy-...
11:26:21 AM	INFO	DROPLET	Staged app 'deploy-service' [Org...
[...]			

The application-specific information displayed in the output is grouped by day (for example, Dec 1. 2017) and typically covers the high-level areas listed and described in the following table:

XS Advanced Application Event Output

Event Information	Description
time	The exact time when the event occurred
level	The severity level assigned to the event, for example, INFO, ERROR, etc.
component	The name of the component that caused the event, for example: APPLICATION, SERVICE, DROPLET, etc.
message	The message text describing the event

Displaying XS Advanced Application Files

If an application is uploaded to the XS advanced run-time platform and staged, the application files are stored in a so called droplet. If an application instances are subsequently started, the files contained in the droplet are

extracted into the file system for the instance to be executed. That means, application files exist on different layers, and the `xs files` command enables you to view the files by choosing the layer you are particularly interested in.

Application Instance Files

To view the files of the currently running application instance, use the `xs files` command with the application name and no further parameters.

Output Code

```
$ xs files deploy-service

Getting files of app "deploy-service" ...

dxwr - META-INF/
dxwr - WEB-INF/
  wr 154 B index.html
dxwr - logs/
```

For each file or directory, `xs files` displays information about the following components:

- The file type and system permissions (d=directory; x=executable; w=writeable; r=readable)
- The size of the file
- The name of the file or directory

To view detailed information about an specific directory tree, add the name of the directory to the command, as shown in the following example.

Output Code

```
$ xs files deploy-service META-INF/

Getting files of app "deploy-service" ...

dxwr - META-INF/
dxwr - WEB-INF/
  wr 154 B index.html
dxwr - logs/
```

To view detailed information about an individual file, specify the full path for the file, as shown in the following example.

Output Code

```
$ xs files deploy-service index.html

Getting files of app "deploy-service" ...

== "index.html" =====
<!DOCTYPE html>
<html>
<head>
<meta charset="ISO-8859-1">
<title>XS2 ALM Service</title>
</head>
<body>Welcome to XS2 ALM Service!
</body>
</html>
```

Application Droplet Files

To view detailed information about the contents of application droplets, use `xs files` command with the `--droplet-index` option, as shown in the following example:

```
$ xs <myAPP> --droplet-files [--droplet-index <index>]
```

The command shows the contents of the root directory of an application's droplet. By default, the latest droplet is displayed. However, you can choose to display details of a specific droplet by specifying its index, for example, with the option `--droplet-index`. For more information about how to display an application's droplet index, use the `xs app <myApp>` command as described in *Displaying Application Details* above.

To download all files associated with an application instance (or droplet) to a target directory on your local disk, use the option `--download <target directory>`, as illustrated in the following example:

```
$ xs <myAPP> --download <target directory>
```

Related Information

[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

[The XS Command-Line Interface \[page 1653\]](#)

12.2.2.1.8 Maintaining Applications in XS Advanced

Deploy applications to the XS advanced run-time environment

The most common task performed when maintaining applications in XS advanced is the deployment of XS advanced applications to the run-time environment. Although XS advanced provides tools to maintain entire business applications consisting of multiple micro-services, the XS command-line interface includes commands that enable administrators to maintain micro-services individually. In addition, the XS CLI also provides commands that help support the maintenance of XS advanced applications during their life cycle:

- [Deploying Multi-Target Applications \(MTA\) \[page 1685\]](#)
- [Deploying individual applications \(microservices\) \[page 1686\]](#)
- [Starting and restarting applications \[page 1687\]](#)
- [Restaging applications \[page 1688\]](#)
- [Configuring Application Health Checks \[page 1688\]](#)

Maintaining SAP Products

The commands `xs install` and `xs uninstall` provided by the Product Installer can be used to maintain SAP products. For more information about using these commands to perform application life-cycle management tasks in XS advanced, see *Related Information*.

Maintaining Multi-Target Applications in XS Advanced

Multi-Target Applications (MTA) are applications composed of one or more microservices. An MTA is delivered in an archive with the file extension `.mtar` (multi-target application archive). The MTAR contains a meta-data description (MTAD, multi-target-application description) as well as the application code itself in a single deployment package. The deployment is performed by the so called Deploy Service, which deploys the individual micro-services specified by an MTAR, creates the corresponding service instances and takes care of connecting the deployed micro-services. To deploy an application in MTAR format, use the command `xs deploy`, as illustrated in the following example:

```
$ xs deploy mybusinessapp.mtar
```

To display a list of all currently deployed MTAs, use the command `xs mtas`, as illustrated in the following example:

Output Code

Listing the Deployed MTAs

```
$ xs mtas
```

```
Getting multi-target apps in org "mymorg" / space "SAP" as XSA_ADMIN...  
Found multi-target apps:
```

mta id	version
com.sap.xsa.admin	1.5.5
com.sap.xs.auditlog.ui	1.0.0
com.sap.xs.jobscheduler	1.6.2
com.sap.core.account	1.1.2
alm-product-installer	1.13.7

To display an overview of the resources and microservices associated with a specific MTA, use the `xs mta` command, as illustrated in the following example:

Output Code

Listing the Deployed MTAs

```
$ xs mta com.sap.xs.jobscheduler
```

```
Getting information for multi-target app "com.sap.xs.jobscheduler"  
in org "orgname" / space "SAP" as XSA_ADMIN...  
Showing information about "com.sap.xs.jobscheduler" version: 1.6.2
```

```
Apps:  
name state inst mem disk urls  
-----  
jobscheduler-backend STARTED 1/1 256MB <unlimited> https://host1:51022  
jobscheduler-broker STARTED 1/1 256MB <unlimited> https://host1:51021  
jobscheduler-dashboard STARTED 1/1 256MB <unlimited> https://host1:51023  
jobscheduler-db STOPPED 0/1 256MB <unlimited> <none>  
jobscheduler-rest STARTED 1/1 256MB <unlimited> https://host1:51017  
jobscheduler-service STARTED 1/1 512MB <unlimited> https://host1:51020  
  
Services:  
name service plan  
-----  
jobscheduler-sbss hana sbss  
jobscheduler-db-container hana hdi-shared
```

jobscheduler-uaa	xsuaa	default
jobscheduler-securestore	hana	securestore

To undeploy all micro-services belonging to a deployed MTA, use the command `xs undeploy`, as illustrated in the following example:

```
$ xs undeploy com.sap.xs.jobscheduler
```

Maintaining Individual Microservices in XS Advanced

To deploy a single XS advanced application, log on the XS advanced run-time instance and use the command `xs push` to deploy the application to the organization and space where you want the application to run, as illustrated in the following example:

```
xs push MY-APP
```

During the deployment the XS advanced runtime performs following tasks:

1. Upload and store the new application files
2. Create or update application meta data
3. Create and bind routes
4. Bind required service instances
5. Choose the appropriate application run time environment
6. Create a droplet of the application
7. Select available execution agents to run the droplet
8. Start the application instances

☰ Output Code

Deploying a Single Application in XS Advanced

```
$ xs push

MY-APP Creating app "MY-APP" in org "orgname" / space "SAP" as XSA_ADMIN...
Creating HTTP route "MY-APP.acme.org" in org "myorg"/space "SAP" as
XSA_ADMIN..

Binding route "https://MY-APP.acme.org" to app "MY-APP"...
Uploading "MY-APP" ...
  Checking which files to upload from /tmp/xs2TestApp ...
  -> "MY-APP" consists of 10 files....
  Uploading 10 new or modified files ...
  Uploading "MY-APP" finished in 265 ms.

Staging app "MY-APP"...
  OUT Detected Java application
  OUT Compiling Java application...
  OUT Java XS Buildpack Version: 1.6.12
  OUT Downloaded 'SAP JVM JRE', version '8.1.33' in 0.559 s.
  OUT Downloaded 'Tomcat Runtime', version '8.5.23' in 0.285 s.
  OUT Downloaded 'XS Authenticator', version '1.6.3' in 0.0 s.
  OUT Downloaded 'SAPJWT', version '1.0.13' in 0.0 s.
  OUT Downloaded 'SAP JVM Memory Calculator', version '1.6.3' in 0.0 s.
Starting app "MY-APP"...
  Starting instances as OS user "sapxsaxsa"
```

```

0 of 1 instances running, 1 starting ...
Showing status and information about "MY-APP"
1 of 1 instances running
  name: MY-APP
  requested state: STARTED
  instances: 1
  memory: 1.00 GB
  disk: <unlimited>
  buildpack: <default>
  urls: https://MY-APP.acme.org

Instances of droplet 1 created at Feb 13, 2018 10:17:36 AM
index  created                state      os user
-----
0      Feb 13, 2018 10:17:49 AM  RUNNING   sapxsaxsa

```

Note

XS advanced applications that use services (for example, database or file-system) are not completely functional the services are provisioned and bound to the application. For more information about maintaining services in XS advanced, see *Related Information*.

Before you push the application to XS advanced, bear in mind the following additional options that might be useful or required when performing a custom deployment:

- Custom host name, port or domain:
You can specify a custom route during deployment. The route is a combination of the domain and subdomain and must be globally unique. This is true whether you specify a portion of the route or allow XS advanced to use defaults.
- Custom start command:
You can specify a custom command to start instances of the application; this custom command replace the automatically detected start command.
- Custom memory limit:
You can specify the maximum amount of memory that each instance of the application can use.

Starting and Restarting XS Advanced Applications

After pushing the application "xs push MY-APP" the XS advanced run-time environment starts the application with the **default** start command of the appropriate build pack. If you want to use a custom start command, use the "-c" option with the `xs push` command, as shown in the following example.

```
$ xs push MY-APP -c "node MY-APP.js"
```

The custom command that you provide with the -c option becomes the default start command and is used for subsequent updates (`xs push MY-APP`) of the application. To reset to the start command used by the default build pack start command, use the special start parameter value "null", as illustrated in the following example:

```
$ xs push MY-APP -c "null"
```

The commands `xs start` and `xs stop` are used to start or stop an application. This always affects all application instances, but it does not change the number of instances configured for scaling the application.

When an application is started, the most recent droplet and the current settings for the application environment variables are used. To restart your application, run the following command:

```
$ xs restart MY-APP
```

Restarting an application stops and restarts it with the most recent droplet and environment variable settings. This is required for “refreshing” the application after any operation that modifies the application's environment, for example, binding a new service or directly setting environment variable values.

i Note

If an environment variable is also consumed by the build pack, then restarting an application might not be enough to effect the desired change. In this case, the application must first be restaged for the change to take effect.

Restaging XS Advanced Applications

Restaging an application compiles a new droplet of the staged application. Staging an application can be useful in the following cases:

- You changed the application environment
- The application must consume an updated application run time
- The application must use an updated set of trust certificates

i Note

A restaged application continues to run during staging; the application only restarts, if you call the `xs restart` command.

To restage your application, use the `xs restage` command, as shown in the following example:

```
$ xs restage MY-APP
```

i Note

Restaging an application compiles a new droplet from the application without updating the application source. To update the application source, push (deploy) the application again.

Configuring XS Advanced Application Health Checks

XS advanced determines the state of application instances by means of so-called application health checks. An application health check is a monitoring process that periodically checks the status of an application instance. If one check fails, the application instance is marked as `CRASHED` and automatically restarted. For more information about the possible status of an XS advanced application instance, for example, `STOPPED`, `STARTING`, or `RUNNING`, etc., see *Displaying Application Information in XS Advanced in Related Information* below.

→ Tip

If tracing is enabled, the health-check status reports are written to the log file of the XS advanced execution agent.

XS advanced supports the following types of application health check:

- **Process:**
The health check examines that the application instance process is running.
- **Port:**
The health check opens a TCP connection on the application instance port and checks that the port is connectable. This is the default health check type.
- **HTTP:**
The health check makes a HTTP GET request to a specified application end point and checks that the response code is 200.

It is possible to set the type of application health check either during application deployment (for example, with the `xs push` command or in the application's `manifest.yml` file) or if the application is already running with the `xs set-health-check` command.

Health Check Configuration Options

You can use the `xs push` command to set the health check for a specific application, using the options listed in the following table:

Application Health-Check Configuration Options

Health-Check Option	Description
<code>-u <HEALTH_CHECK_TYPE></code>	The application health check type. Valid values are port, process or http. The default is port.
<code>--health-check-timeout <TIMEOUT></code>	The XS advanced runtime will wait at most the specified time for a successful health check after the initial startup of an instance. If no health check was successful after the timeout was reached, the instance will be marked as crashed and the instance will be terminated. The default is to wait indefinitely.
<code>--health-check-http-endpoint <HEALTH_CHECK_ENDPOINT></code>	The end point for the HTTP request provided by an application. This option is only relevant for the health-check type http. The default value is <code>"/</code> .
<code>--invocation-timeout <HTTP_INVOCATION_TIMEOUT></code>	The timeout for each HTTP GET request. This option is only relevant for the health-check type http. The default value is 1 second.

In the following example, the command pushes an application and configures the health check type 'http' with a health check timeout of 180 seconds using the HTTP endpoint `/healthcheck` and a GET request timeout of 5 seconds:

```
$ xs push demo-app -u http --health-check-timeout 180 --health-check-http-endpoint /healthcheck --invocation-timeout 5
```

The following example shows how to specify the same configuration in the application manifest:

```
---
```

```
...
health-check-type: http
health-check-http-endpoint: /healthcheck
invocation-timeout: 5
```

If the application has already been pushed, you can use the `xs set-health-check` command to configure health checks, as illustrated in the following example:

```
xs set-health-check <application> http --endpoint /healthcheck --invocation-
timeout 5
```

Note

You must restart the application to enable the configured health checks.

Viewing the Health Check Configuration for an Application

You can use the `xs get-health-check` display the current health-check configuration, as illustrated in the following example:

```
$ xs get-health-check demo-app

name :                               demo-app
health check type :                   http
endpoint :                             /
invocation timeout (in seconds) :     1

OK
```

Related Information

[Displaying Application Information in XS Advanced \[page 1676\]](#)

[Installing and Updating Products and Software Components in SAP HANA XS Advanced Model \[page 960\]](#)

[Maintaining Services in XS Advanced \[page 1692\]](#)

[Maintaining the XS Advanced Application Environment \[page 1702\]](#)

[Maintaining Trust Certificates in XS Advanced \[page 1706\]](#)

12.2.2.1.9 Scaling Applications in XS Advanced

Configure applications to handle increased traffic on demand.

Scaling an application can enable it to handle increased traffic on demand, for example, to meet the needs associated with increased user load. The command `xs scale` can be used to increase or decrease the number of running application instances (also known as horizontal scaling) or to change the resource limit of existing application instances (vertical scaling).

Horizontal Scaling

You can horizontally scale your application by specifying the number of application instances that should be allowed, as illustrated in the following example, which sets the maximum number of instances to five (5):

```
$ xs scale MY-APP -i 5
```

This configures the XS advanced run-time environment to increase or decrease the number of instances of the specified application to match the number of instances specified.

i Note

The consumption of file-system and memory resources increases with the number of application instances. In addition, the application itself determines whether it supports horizontal scaling, that is, whether it can manage several application instances running in parallel.

A round-robin policy is used to distribute newly created application instances among available hosts bearing the role `xs_worker`. You can use host pinning to refine the set of `xs_worker` hosts where instances of specific applications will be scheduled to run. For more information about host pinning, see *Related Information*. To decrease the number of application instances, the same round-robin policy is used to remove application instances from affected `xs_worker` hosts. The load generated by requests sent to a scaled application is automatically balanced across all application instances by the XS advanced Platform Router. Although, by default, a weighted round-robin policy is used to balance the load associated with requests among application instances, it is nonetheless possible to choose from the other load-balancing algorithms provided by the SAP WebDispatcher.

Vertical Scaling

To change the memory limits for existing applications, use the command "xs scale" with the option "-m", for example:

```
$ xs scale MY-APP -m 1G
```

In this example, the memory limit of the application `MY-APP` is changed to 1 GB. To define the unit of size, you can either use "M" (for megabytes) or "G" for gigabytes. For the new settings to be effective, restage and restart the corresponding application, for example, with the commands `xs restage` and `xs restart`, respectively. For more information about restaging and restarting application in XS advanced, see *Related Information*.

! Restriction

Memory limits currently are only valid for Java applications.

Related Information

[Maintaining Host Pinning \[page 1717\]](#)

[Maintaining Applications in XS Advanced \[page 1684\]](#)

12.2.2.110 Maintaining Services in XS Advanced

Use the `xs` command-line interface to maintain so-called “Backing Services” for applications running in the XS advanced run-time platform.

Backing Services are the medium by which applications running in the XS advanced run-time platform can access resources, for example: a database, an offering for audit logging, or user authentication.

Backing Services are connected with applications by creating so called service instances within XS advanced. During the creation of a service instance, XS advanced creates the backing-service resource by means of a so-called “service broker”. By adding custom service brokers, arbitrary backing services can be connected to XS advanced. At some point, the service instance is bound to an application. In this way, all the information required for access to the backing service is injected into the process environment of an application instance.

- [Listing Backing Services in the Market Place \[page 1692\]](#)
- [Maintaining Service Instances \[page 1693\]](#)
- [Creating User Provided Services \[page 1694\]](#)
- [Binding Service Instances to Applications \[page 1695\]](#)
- [Maintaining Service Keys \[page 1695\]](#)
- [Maintaining Service Brokers \[page 1696\]](#)
- [Maintaining Syslog Drain Services \[page 1697\]](#)

Listing Backing Services in the Market Place

Typically the service broker provides credentials to the application to access the service. In XS advanced, the list of available backing services provided by service brokers is displayed in the service “market place”, as shown in the following example:

```
⌵ Output Code

$ xs marketplace

Getting services from marketplace...

service      plans      description
-----
fs-storage   free       xs file service provides.
xsuaa        default, devuser, space  Manage app authorizati...
hana         hdi-shared, sbss, schema, securestore  SAP HANA database
managed-hana hdi-shared, schema, securestore  Creates service instan...
auditlog     free       Audit log broker on XSA..
```

The output from the `xs marketplace` command displays the following information:

- `service`
The name of the backing service
- `plan`
The available service plans a particular backing service provides. A service plan determines the type or category of service that an individual backing service provides. For example the `hana` service can be used to get a single container (`schema`) or provide access to the SAP HANA Secure Store (`securestore`).

- `description`
A short summary of the backing service

Maintaining Service Instances

To make use of a service, an instance of the service must be created and an the XS advanced application must be bound to the specified service instance. When a service instance is created, resources are allocated within the respective backing service. To create a service instance, use the command `xs create-service` and choose a service name and a service plan from the marketplace, as illustrated in the following example:

```
$ xs create-service hana hdi-shared myservice  
Creating service "myservice"  
OK
```

You can configure service instances by passing a set of parameters with the option `-c` in JSON format. You can also add tags to the service instance with the option `-t`; the information is passed to the bound applications. Service instances are created within a particular space and cannot be used from other spaces.

The `xs update-service` command enables you to update a service instance, for example, by modifying its service plan as well as any parameters or tags:

```
$ xs update-service myservice -t mytag  
Updating service instance "myservice"...  
OK
```

⚠ Caution

The corresponding service broker determines what happens during a service update. For example, changing the plan of an existing service could result in a new allocation of service resources within the backing service, and this could lead to the deletion of all previous data represented by this service instance.

The `xs rename-service` command enables you to rename a service instance, as illustrated in the following example:

```
$ xs rename-service myservice newservice  
OK
```

i Note

Renaming a service instance has no effect on the backing resource represented by the service instance.

The `xs delete-service` command enables you to remove a service instance along with all associated data, as illustrated in the following example:

```
$ xs delete-service myservice  
Really delete service instance "myservice"? (y/n) > y  
Deleting service instance "myservice"...  
OK
```

Since a service instance is maintained and persisted not only at the service broker but also at the XS Controller, the service-related information should be consistent in both places. If the service information is lost at the

service broker, for example, because it has been manually deleted, you can use the purge option (`xs delete-service --purge`) to forcibly remove the service information from the XS Controller as well.

To display a list of all currently available service instances in a space, use the command `xs services` as illustrated in the following example:

☰ Output Code

Viewing Available Service Instances

```
$ xs services

Getting services from marketplace...

name                service    plan          bound apps
-----
auditlog-db-container hana      hdi-shared    auditlog-db, auditlog-...
auditlog-sbss        hana      sbss          auditlog-server, auditlog
deploy-service-auditlog auditlog  free          deploy-service
deploy-service-fss   fs-storage free          deploy-service
deploy-service-ss    hana      securestore  deploy-service
deploy-service-database hana     schema       deploy-service
deploy-service-uaa   xsuaa    default      deploy-service
product-installer-dbase hana     schema       product-installer
component-registry-dbase hana     hdi-shared    component-registry-db,...
```

Creating User Provided Services

Since user-provided services do not require a service broker they are not chosen from the service catalog. Instead, the credentials for a user-provided services are provided when creating the service instance, as illustrated in the following example:

☰ Output Code

User Credentials in JSON for a User-Provided Service

```
$ xs create-user-provided-service my-up-service -p
'{"host":"example.org","username":"admin","password":"pa55woRD"}'

Created environment (excerpt):
{
  "name" : "my-up-service",
  "credentials" : {
    "password" : "pa55woRD",
    "host" : "example.org",
    "username" : "admin"
  }
}
```

→ Tip

The combination of operating system, shell, or terminal type determines the “quote” or “escape” characters required when providing the parameter string in JSON format. For more information about the correct quoting, see `xs help create-user-provided-service`.

The `xs` CLI also has an interactive mode, which you can use to specify the fields of the service credentials as parameters, for example, with the `-p` option and a comma-separated list.

Output Code

Interactive User Credentials for a User-Provided Service

```
$ xs create-user-provided-service my-up-service -p host,username,password
host> example.org
username> admin
password> pa55woRD

Created environment (excerpt):
{
  "name" : "my-up-service",
  "credentials" : {
    "password" : "pa55woRD",
    "host" : "example.org",
    "username" : "admin"
  }
}
```

Binding Service Instances to Applications

To connect an application with a particular backing service resource, you must bind the corresponding service instance to the application using the `xs bind-service` command, as illustrated in the following example:

```
$ xs bind-service myapp myservice
OK
TIP: Use 'xs restart' to ensure your env variable changes take effect
```

It is also possible to pass binding parameters in JSON format by using the option `-c`. The effect of binding parameters depends very much on the corresponding backing service. After restarting the bound application, credentials required for to access to the backing service can be found in the environment variable `<VCAP_SERVICES>` in the application environment. To remove a service-binding, call the command `xs unbind-service`, as illustrated in the following example:

```
$ xs unbind-service myapp myservice
OK
```

Maintaining Service Keys

Service keys provide a way to query the credentials required for a service without having to bind the service instance to an application. To create a service key use the `xs create-service-key` command, as shown in the following example:

```
$ xs create-service-key myservice myservicekey

Creating service key "myservicekey" for service instance "myservice" ...
OK
```

Service keys are created within a particular space and cannot be used from other spaces. After creating a service key, you can view the service credentials by calling the command `xs service-key`, as illustrated in the following example:

Output Code

```
$ xs service-key my-service my-service-key

Getting service key "my-service-key" for service instance "my-service" ...
{
  "host" : "host1.acme.org",
  "user" : "1EBB56E88DYS76",
  "password" : "Dd2f8RVCP8gr1[...]v"
}
OK
```

To display a list of all existing service keys, use the command `xs service-keys`:

Output Code

```
$ xs service-keys

Getting service keys in org "orgname" / space "SAP" as XSA_ADMIN...

service instance  name
-----
myservice         myservicekey
```

To remove a service key, use the command `xs delete-service-key`:

Output Code

```
$ xs delete-service-key myservice myservicekey

Really delete service key "myservicekey" for service instance "myservice"?
(y/n) > y
Deleting service key "myservicekey" for service instance "myservice" ...
OK
```

Maintaining Service Brokers

The service broker interface provides a way to connect arbitrary backing service with the XS advanced platform. XS advanced implements the open service broker API. To view the service brokers currently connected to the XS advanced platform, use the `xs service-brokers` command, as shown in the following example:

Output Code

```
$ xs service-brokers

Getting service brokers...
Found service brokers:

name          url
```

```
-----
fs-storage          https://acme.org:30033/v2/fs-service
uaa-security        https://acme.org:30033/uaa-security
hdi-broker          https://acme.org:30033/hdi-broker
instance-manager    https://acme.org:30033/instance-manager
auditlog            https://acme.org:30033
```

You can add your own service broker with the command `xs create-service-broker`, as illustrated in the following example:

Output Code

```
$ xs create-service-broker mybroker user passwd https://acme.org:8080/activemq
OK
```

After a service broker has been created, the backing services provided by this service broker immediately show up in the service catalog, also known as the marketplace.

To rename an existing service broker, use the following command:

Output Code

```
$ xs rename-service-broker mybroker newbroker
OK
```

If a service broker is moved to a different URL or the credentials for a service broker change, you can update an existing service broker to reflect the new URL and new credentials by calling the command `xs update-service-broker`, as illustrated in the following example:

Output Code

```
$ xs update-service-broker mybroker user passwd https://acme.org:1080/activemq
OK
```

To remove an existing service broker from the XS advanced platform, use the following command:

Output Code

```
$ xs delete-service-broker mybroker

Really delete service broker "mybroker"? (y/n) > y
Deleting service broker "mybroker"...
OK
```

Maintaining Syslog Drain Services

In order to stream application logs to an external log management service, for example an Elastic (ELK) stack, you can make use of so called "syslog-drain services". These services define a URL to the `syslog-compatible`

end point of the log-management service: the so-called syslog drain. The same application logs that can be streamed to a `syslog` drain can be accessed with the command `xs logs`, as described in *Displaying Application Information* in *Related Information*. XS advanced supports syslog drains which are able to parse log messages according to the standards described in RFC 5424.

To start streaming logs from your application to the syslog drain, you need to create a user-provided service and bind it to your application. The following example shows how to use the `xs` command-line interface to create the corresponding user-provided service; you need to specify a service name and, by means of the parameter `-l`, the URL of the syslog drain:

```
$ xs create-user-provided-service my-drain -l syslog://syslog-drain.example.org:1234
OK
```

The log messages can be delivered via TCP, TCP over TLS, or HTTPs, as shown in the following table, which provides an overview of the format of the syslog-drain URL that XS advanced expects for the supported protocols:

URL Formats for Syslog Drain in XS Advanced

Protocol	Syslog-drain URL Format
TCP	<code>syslog://syslog-drain.example.org:1234</code>
TCP over TLS	<code>syslog-tls://syslog-drain.example.org:1234</code>
HTTPs	<code>https://syslog-drain.example.org:1234</code>

Next, you need to bind the created service to your XS advanced application, as illustrated in the following example:

```
$ xs bind-service my-app my-drain
OK
TIP: Use 'xs restart' to ensure your env variable changes take effect
```

After a short delay all new logs generated by the bound application will be streamed to the syslog drain. A restart of the application is not required, as the syslog-drain service binding does not affect the environment variables of the application.

If you are streaming logs to a syslog drain that uses TLS, make sure that the required trust certificates are available to the XS Controller. For more information about how to configure trust certificates please see *Maintaining Trust Certificates* in *Related Information* below.

Bear in mind that additional configuration might be required within the log-management service of your choice. For more information, see *The Syslog Protocol* in *Related Information*.

Related Information

[XS Advanced Platform Components \[page 1648\]](#)

[Displaying Application Information in XS Advanced \[page 1676\]](#)

[The Syslog Protocol \(ietf.org\) !\[\]\(2756f1a8aaf2ed04e6d077102c4b8691_img.jpg\)](#)

[Maintaining Trust Certificates in XS Advanced \[page 1706\]](#)

12.2.2.11 Maintaining Application Routes in XS Advanced

Use the `xs` command-line interface to maintain routes for XS advanced applications.

Routes provide applications and their instances with a public HTTP endpoint. A route represents a URL that can be mapped to one or more applications. HTTP traffic that is sent to a route is forwarded by the Platform Router to one of the instances of the mapped applications. The selection of the application instance is decided by a round-robin algorithm. The Platform Router acts as the central routing and load-balancing component within XS advanced.

Usually a route is mapped to a single application. When scaling the application to multiple instances, traffic sent to the route is load-balanced between all instances. Mapping a route to multiple applications might for example make sense in cases where a blue-green deployment is performed.

Routing Modes

During installation of SAP HANA XS advanced model you can select between two routing modes: **ports** and **host names**. If the administrator configures the routing mode “ports”, the XS advanced platform creates routes based on distinct **ports**, as illustrated in the following URLs representing different port-based routes:

- `https://acme.org:50000`
- `https://acme.org:50001`

If the administrator configures the routing mode “host names”, the XS advanced platform creates routes based on distinct sub-domains, as illustrated in the following URLs representing different host-name-based routes:

- `https://app1.acme.org`
- `https://app2.acme.org`

→ Tip

For more information about technical prerequisites and security implications of the two routing modes see SAP Note [2245631](#), *Domains and Routing Configuration for SAP HANA Extended Application Services, Advanced Model*.

In the “hostnames” routing mode, the XS advanced platform can also create routes based on ports, which is especially useful if you also want to use TCP routes, which always require a dedicated port.

To simplify the creation of port-based routes, you can configure a default port range for the platform, for example, by setting XS advanced platform parameters. In this way, a free port from the declared port range can be used if no port is explicitly defined when creating a route. The port range should be treated as **reserved**, and SAP HANA XS advanced expects that no other processes on the SAP HANA hosts use ports within the declared port range. If you are running multiple SAP HANA XS advanced systems on the same host, it is important to separate the default port range for each XS advanced system from each other.

→ Tip

For more information about setting a default port range for XS advanced, see SAP Note [2507070](#), *Multiple XS Advanced Systems on the Same Host*.

Viewing Routes

To see which routes have been mapped to XS advanced applications in the current space, use the `xs routes` command, as illustrated in the following example:

Output Code

Route List in "hostnames" Routing Mode

```
$ xs routes

Getting routes in org "XSA" / space "SAP" as XSA_ADMIN...

host                domain    port    path  type  apps
-----
auditlog-server     acme.org          /    HTTP  auditlog-server
auditlog-broker     acme.org          /    HTTP  auditlog-broker
deploy-service      acme.org          /    HTTP  deploy-service
xsa-sap-product-installer acme.org          /    HTTP  product-installer
xsa-sap-auditlog-odata acme.org          /    HTTP  auditlog-odata
xsa-sap-auditlog-ui acme.org          /    HTTP  auditlog-ui
                    acme.org  50510  /    TCP   hello-world-tcp
                    acme.org  50500  /    TCPS  hello-world-tcps
```

The information displayed in the example above includes the following details:

XS Advanced Application Details

Property	Description
host	The host of the specified route. The host can only be set where hostname-based routing is used.
domain	The domain name used by the specified route. → Tip For more information about domains, see <i>Maintaining Domains</i> in <i>Related Information</i> .
port	The port that is dedicated to the specified route. The port can only be set if the route is a port-based route. → Tip Host-name-based routes use the default router port.
path	The URL path of the specified route. Using different URL paths you can create multiple routes with the same host or port.
type	The protocol used for the specified route, which can be one of the following: HTTP, TCP, or TCPS.
apps	The applications bound to the specified route

Creating and Mapping Routes

You can use the `xs map-route` and `xs unmap-route` commands to map routes to applications and remove the mappings between applications and routes. Similarly, the commands `xs create-route` and `xs delete-route` can be used to create new routes and delete them.

You can also define the URL path when creating a route. The URL path defaults to "/", but using different URL paths enables you to create multiple routes that use the same sub-domain or port, as illustrated in the following examples:

- `https://app.acme.org/frontend`
- `https://app.acme.org/backend`

The default type of a route is HTTP. This means that the traffic on this route must adhere to the HTTPs (HTTP over SSL) protocol, which is used within XS advanced by default. It is also possible to change this configuration globally, in order to use plain HTTP routes. If a route is created with type HTTP, the XS advanced Platform Router can enable certain HTTP features for this route, for example: Sticky Sessions, HTTP access logs, or certain timeouts. The XS advanced platform is also able to provide additional information to the application instances through HTTP headers (for example, `X-Forwarded-Proto` and `X-Forwarded-Port`). Routes, which are based on hosts, can only be used with the HTTP type, as the Platform Router routes these requests to the correct application instances based on the `HTTP Host` header.

It is also possible to create routes of type "TCP" or "TCPs". It is important to remember that the XS advanced Platform Router does not expect a specific protocol on TCP routes; it simply forwards the TCP traffic to the application instances using the mapped route. TCP routes must, however, use their own distinct port and are not allowed to use hosts. All HTTP features are disabled for TCP routes, so it only makes sense to use these routes if the application requires a protocol other than HTTP. In addition, since the XS advanced Platform Router does not terminate SSL, for TCP routes you can also terminate SSL traffic at the application instance.

You can also use the route type "TCPs", which expects at least SSL-encrypted TCP traffic and allows the XS advanced Platform Router to terminate the SSL connection. The application instance receives plain TCP traffic. The advantage of using "TCPs" is that the XS advanced application does not require any SSL configuration, and you can continue to rely on the domain certificates you have configured within SAP HANA XS advanced.

Related Information

[Maintaining Domains in XS Advanced \[page 1708\]](#)

[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

[The XS Command-Line Interface \[page 1653\]](#)

12.2.2.112 Maintaining the XS Advanced Application Environment

Use environment variables to configure an XS advanced application's behavior.

When an application instance is started in XS advanced, several configuration parameters are injected into the application's process environment; the parameters specify information about the following elements:

- Application properties
- Information about backing services
- Custom environment values

→ Tip

You can use environment variables to transport custom configuration parameters to application instances.

Viewing the Application Environment

The command `xs env` displays an overview of an application's environment. The command output is split into two sections: `System-Provided` and `User-Provided`, as illustrated in the following example:

```
$ xs env demo-app

Getting env variables for app "demo-app"...
OK

System-Provided:
{
  "VCAP_APPLICATION" : {
    "start" : "2018-01-02 16:51:41 +0100",
    "application_id" : "927c0822-3734-4c91-9796-432ddd1a9ed4",
    "instance_id" : "927c0822-3734-4c91-9796-432ddd1a9ed4",
    "space_id" : "eccc5bd5-934b-4472-b32c-009d6497e93b",
    "application_name" : "demo-app",
    "organization_name" : "myorg",
    "space_name" : "PROD",
    "started_at_timestamp" : "1514908301002",
    "started_at" : "2018-01-02 16:51:41 +0100",
    "state_timestamp" : "1513239094881",
    "full_application_uris" : [ "https://demo-app.example.org" ],
    "application_uris" : [ "demo-app.example.org" ],
    "uris" : [ "demo-app.example.org" ],
    "version" : "3d397398-0f1a-4d07-b9e6-88f33d8387b0",
    "application_version" : "3d397398-0f1a-4d07-b9e6-88f33d8387b0"
  },
  "VCAP_SERVICES" : {
    "hana" : [ {
      "name" : "demo-app-database",
      "label" : "hana",
      "tags" : [ "hana", "database", "relational" ],
      "plan" : "schema",
      "credentials" : {
        "schema" : "USR_EHLDI8B567MXSPX1BLGSJJW7J",
        "password" : "1h0FH1R9o7WvC5AHmI_522DU8kLPIYKM",
        "driver" : "com.sap.db.jdbc.Driver",
        "port" : "31313",
        "host" : "962.acme.com",
```

```

    "db_hosts" : [ {
      "port" : 31313,
      "host" : "962.acme.com"
    } ],
    "user" : "USR_EHLDI8B567MXSPX1BLGSJJW7J",
    "url" : "jdbc:sap://962.acme.com:31313/?
currentschema=USR_EHLDI8B567MXSPX1BLGSJJW7J"
  }
}
User-Provided:
  custom-value: 1234

```

System-Provided Environment Variables

System-provided environment variables are calculated by XS advanced and cannot be changed the user. The following environment variables contain application and service-binding properties:

- `VCAP_APPLICATION`
Contains application properties in JSON format, for example, the time stamp of the last application start time, unique application identifiers, organization and space coordinates, and application URLs.
- `VCAP_SERVICES`
Contains information about the bound backing services in JSON format. For each backing service, the information includes the credentials required for access to the specified backing service.

For more details about *Maintaining Services*, see *Related Information* below.

User-Provided Environment Variables

Custom user-provided environment variables can be set by the user or are set by the deployer of the application.

Modifying the Application Environment

To add a new environment variable or alter the value of an existing environment variable, use the command `xs set-env`, as shown in the following example:

Note

How an application consumes the value you set with the `xs set-env` command depends very much on the application itself.

Output Code

Setting Environment Variables in XS Advanced

```
$ xs set-env demo-app "http_proxy" "proxy:8080"
```

```
Setting env variable "http_proxy" to "proxy:8080" for app "demo-app"...
```

```
OK
```

```
TIP: Use 'xs restage' followed by 'xs restart' to ensure your env variable
changes take effect
```

To remove a new environment variable from the existing application environment, use the command `xs unset-env`, as shown in the following example:

☰ Output Code

Setting Environment Variables in XS Advanced

```
$ xs unset-env demo-app http_proxy

Removing variable "http_proxy" from environment of app "demo-app"...
OK
TIP: Use 'xs restage' followed by 'xs restart' to ensure your env variable
changes take effect
```

Any change to the application environment is visible in the output of the `xs env` command immediately. However, the effect of the change will only be applied to the application after the application has been restarted.

→ Tip

If the buildpack that compiles the application consumes the altered environment variables, you must restage the application, too.

Global Environment Variables

It is possible to set global environment values for all applications; this avoids having to set the environment for each application individually. A common use case for such a scenario would be when applying proxy settings.

Running Environment Variable Groups

To set a group of environment variables for all newly started applications, use the command `xs set-running-environment-variable-group`. You can pass a group of variables by specifying a JSON structure as a command-line option, as illustrated in the following example:

i Note

This command overwrites any existing running environment variable group you have configured.

☰ Output Code

Setting Running Environment Variables in XS Advanced

```
$ xs set-running-environment-variable-group '{"no_proxy":"localhost"}'

Setting the contents of the running environment variable group as XSA_ADMIN...

Variable Name      Assigned Value
-----
no_proxy           localhost
```

To view the currently configured running environment variable group, use the command `xs running-environment-variable-group`, as shown in the following example:

i Note

Environment variables in the running environment variable group have no effect on the staging processes.

≡ Output Code

Displaying the Running Environment Variables Group in XS Advanced

```
$ xs running-environment-variable-group

Retrieving the contents of the running environment variable group as
XSA_ADMIN...

Variable Name      Assigned Value
-----
no_proxy           localhost
```

Staging Environment Variable Groups

To set a group of environment variables that is only exposed to build packs during the application staging process, use the command `xs set-staging-environment-variable-group`, as illustrated in the following example:

i Note

This command overwrites any existing **staging** environment variable group you have configured.

≡ Output Code

Setting Staging Environment Variables in XS Advanced

```
$ xs set-staging-environment-variable-group '{"no_proxy":"localhost"}'

Setting the contents of the staging environment variable group as XSA_ADMIN...

Variable Name      Assigned Value
-----
no_proxy           localhost
```

To view the currently configured staging environment variable group, use the command `xs staging-environment-variable-group`, as shown in the following example:

≡ Output Code

Displaying the Staging Environment Variables Group in XS Advanced

```
$ xs staging-environment-variable-group

Retrieving the contents of the staging environment variable group as
XSA_ADMIN...

Variable Name      Assigned Value
-----
no_proxy           localhost
```

Related Information

[Maintaining Services in XS Advanced \[page 1692\]](#)

12.2.2.1.13 Maintaining Trust Certificates in XS Advanced

Use the `xs` command-line interface to maintaining trust certificates for the XS advanced run-time platform.

The XS advanced run-time environment and the applications running in XS advanced use SSL connections wherever possible. Applications connect to each other using HTTPs and can also be configured to use SSL when connecting to the SAP HANA database. Although the XS advanced platform components (for example, the XS Controller) use the SAP HANA system PKI to securely communicate with each other and with the SAP HANA database, nonetheless it might be necessary to configure additional SSL trust certificates, especially when connecting securely to end points that are located outside the XS advanced run-time environment. Since application run times (for example, Java or Node.js) usually already include their own certificate trust stores, configuring additional SSL trust certificates might only be necessary when using certificates, which are not signed by a well known certificate authority (CA). This usually is the case for certificates signed by a corporate CA or self-signed certificates.

The XS advanced run-time platform includes tools that enable administrators to maintain and manage the certificates used to set up trusted relationships; you can use the tools to upload custom trust certificates to the platform. The trust certificates are used by the platform components (for example, the XS Controller) and can be propagated to the applications by the build packs. For example, the propagation of the certificates is performed by default when using either the SAP Node.js or the SAP Java build pack. However, you can use **custom** trust certificates, for example, if you have configured a custom “certificate authority” within your corporation.

By adding the root certificate of your certificate authority, you enable any application to establish fully trusted SSL sessions with corporate servers, even those configured outside the SAP HANA system. In addition to custom trust certificates, the XS advanced run-time platform can be used to manage **domain** certificates. For more information, see *Maintaining Domains* in *Related Information* below.

→ Tip

Domain certificates are treated the same way as custom trust certificates; both types of certificates are available to **all**.

Viewing Trust Certificates

To display a list of custom trust certificates, use the `xs trusted-certificates`, as illustrated in the following example:

≡ Output Code

```
$ xs trusted-certificates
```

```
Retrieving the list of trusted certificates as XSA_ADMIN...
```

```

Alias: CORPORATE_CA
-----
Subject:          CN=CORPORATE Root CA,O=SAP SE,L=Walldorf,C=DE
Issuer:           CN=CORPORATE Root CA,O=SAP SE,L=Walldorf,C=DE
Valid from:       Sun Jan 01 12:00:00 UTC 2017
Valid until:      Thu Jan 01 11:59:59 UTC 2032
Signature algorithm: SHA256withRSA

Alias: HANA_SSL
-----
Used within HANA Broker service bindings
Subject:          CN=hanahost,O=SAP SE,L=Walldorf,C=DE
Issuer:           CN=hanahost,O=SAP SE,L=Walldorf,C=DE
Valid from:       Sun Jan 01 12:00:00 UTC 2017
Valid until:      Mon Jan 01 11:59:59 UTC 2018
Signature algorithm: SHA256withRSA
Alias: SSO_CA
-----
Used for validation of certificate-based client authentication
Subject:          CN=CORPORATE SSO CA,O=SAP SE,L=Walldorf,C=DE
Issuer:           CN=CORPORATE SSO CA,O=SAP SE,L=Walldorf,C=DE
Valid from:       Sun Jan 01 12:00:00 UTC 2017
Valid until:      Thu Jan 01 11:59:59 UTC 2032
Signature algorithm: SHA256withRSA

```

The `xs trusted-certificates` command displays a list of all trust certificates including the alias for a certificated (if defined) and some additional information that is extracted from the certificate itself, for example, the subject, the issuer and the period for which the certificate is valid.

Adding Trust Certificates

To upload a custom trust certificate, use the command `xs trust-certificate`, as illustrated in the following example:

Output Code

```
$ xs trust-certificate <ALIAS> -c <PATH>
```

In the example command output above, `<ALIAS>` is an arbitrary name and has no special meaning. `<PATH>`, however, must point to a valid certificate file, which must be a X.509 certificate in the PEM format.

⚠ Caution

All trust certificates are passed to all applications and can be used to validate SSL sessions to any server.

The core components of the XS advanced platform, for example, the XS Controller and the XS UAA, rely on these trust certificates to establish trusted SSL connections.

SAP HANA SSL Trust Certificates

If you want to use SSL to encrypt connections between the SAP HANA database and SAP HANA XS advanced, you must upload to the XS advanced trusted-certificate store the root certificate of the server certificate, which

is used by your SAP HANA database. When adding this trusted certificate, pass the additional option "--hana-broker". This is required to enable the propagation of the certificate to applications, which use the SAP HANA Broker to obtain details of their SAP HANA database connection. For example, Node.js applications rely on this information to determine which certificate should be used to validate the database SSL connection.

Only one trust certificate can be used with the SAP HANA Broker. If you are using multiple tenant databases, all server certificates of these tenant databases must be signed by the same root certificate. If you are adding trust certificates to remote SAP HANA systems that are not managed by the SAP HANA Broker (for example, SAP HANA systems managed within SAP HANA Cockpit 2.0), simply upload these certificates as standard trust certificates.

→ Tip

For information how to setup SSL encryption between SAP HANA XS advanced model and the SAP HANA database it is running on, see SAP Note [2300943](#) (Enabling SSL encryption for database connections for SAP HANA extended application services, advanced model).

Client Authentication Trust Certificates

The XS advanced platform router supports the validation of client certificates to allow certificate-based authentication. The platform router also has the option to request client certificates if at least one trust certificate for client authentication has been configured in XS advanced. The validated client certificate is passed on to platform services and applications in the HTTP header "X-Forwarded-Client-Cert".

The certificates used by the platform router to validate a client certificate during the SSL handshake must be uploaded to XS advanced using the `xs trust-certificate` command together with the `--client-auth` option. Certificates that are uploaded using the `--client-auth` option are not propagated to applications.

i Note

It is strongly recommended to use root CA certificates as client-authentication trust certificates; it is not recommended to upload user-specific client-authentication trust certificates.

Related Information

[Maintaining Domains in XS Advanced \[page 1708\]](#)

[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

[The XS Command-Line Interface \[page 1653\]](#)

12.2.2.14 Maintaining Domains in XS Advanced

Use the `xs` command-line interface to maintain domains and domain certificates for the XS advanced run-time platform.

During installation you need to specify a default domain. This domain is used when creating URLs for the platform components and when creating the routes of the applications created by the initial installation. The default domain is a shared domain, which means it can be used from within any organization. When creating additional domains one can limit the usage of a domain to a specific organization. The domain are called private domains and can then only be used by routes created in that organization. The example URLs show what two routes using different domains might look like: Domains are the building blocks for application routes. Routes reference domains and specify a subdomain or a port on a domain, which can then be bound to applications. You can have multiple domains in your XS advanced system. DNS entries for these domains need to point to the Platform Router.

- `https://app.acme.org`
- `https://app.myotherdomain.com`

To display a list of all domains, run the following command:

```

Sample Code

xs domains

Getting domains...

name                type                org
-----
acme.org             shared (default)
myotherdomain.com   shared
privatedomain.com   private             XSA
  
```

The information displayed in the example above includes the following details:

XS Advanced Application Details

Property	Description
shared (default)	Domains are the building blocks for application routes. Routes reference domains and specify aThe default domain that is available to (shared between) all organizations
shared	A domain that is available to all organizations
private	Domains that are only available to a single organization

If XS advanced is configured to use the host-name-based routing mode, the XS advanced platform automatically creates subdomains for a domain when a new route is created. To ensure that this is possible, DNS entries for the domains need to include a wild-card DNS entry, which needs to point to the XS advanced Platform Router. For more information about the prerequisites for configuring and using domains in XS advanced, see SAP Note [2245631](#), *Domains and routing configuration for SAP HANA extended application services, advanced model*.

Managing Domain Certificates

You can configure SSL certificates for every domain: the SSL certificate for a domain is used for every application route that uses this domain. By default, the XS advanced platform automatically generates a self-signed SSL certificate for every domain; the XS advanced administrator can replace this self-signed SSL

certificate with a certificate signed by a well-known certificate authority. If the “host names” routing mode is enabled, it is necessary to deploy a wild-card certificate to each domain, since routes for a domain are created as subdomains. If you are using the “ports” routing mode, wild-card certificates are not needed.

i Note

Wild-card certificates include the domain name (for example, "example.org") and an additional wildcard entry such as "*.example.org" in the certificate, as illustrated in the following example.

To display a list of domain certificates for each configured domain, use the `xs domain-certificates`, as illustrated in the following example:

Output Code

```
$ xs domain-certificates

Retrieving the list of domain certificates as XSA_ADMIN...

Domain: acme.org
-----
Last updated at:           Sun Jan 01 12:00:00 UTC 2017
Created by:                Platform (self-signed)
Subject:                   CN=acme.org,OU=XS,O=ACME SE,C=GH
Subject alternative names: [[2, .acme.org], [1, *.acme.org]]
Issuer:                    CN=acme.org,OU=XS,O=ACME SE,C=GH
Valid from:                Sun Jan 01 12:00:00 UTC 2017
Valid until:               Sun Jan 01 12:00:00 UTC 2018
Signature algorithm:       SHA256withRSA
```

In addition to domain certificates, the XS advanced run-time platform can be used to manage custom trust certificates. For more information, see *Maintaining Trust Certificates* in *Related Information* below.

→ Tip

Domain certificates are treated the same way as custom trust certificates; both types of certificates are available to **all** XS advanced applications.

For more information about how to set certificates for the XS advanced run-time platform, see SAP Note [2243019](#).

Setting Domain Certificates in PEM Format

To configure your own certificate for a domain you can use the "xs set-certificate" command. The command requires a private key in PKCS8 PEM format and the full certificate chain in X.509 PEM format. The following example walks you through the process of creating a CA-signed certificate and setting it as the domain certificate. It uses the command-line tool `openssl`.

→ Tip

For more information about creating CA-signed certificates, see SAP Note [2243019](#) (*Providing SSL certificates for domains defined in SAP HANA extended application services, advanced model*).

As part of the process of creating a domain certificate, you must first create a certificate-signing request, which you would normally define in the `openssl` configuration files, but can also do on the command line, in a bash

script, as illustrated in the following very simple example, in which the contents of the `subjectAltName` and `Subject` fields must be adapted to fit to the corresponding domain name:

```
openssl req -new -sha256 -newkey rsa:4096 -nodes \  
-keyout domain.key \  
-subj "/CN=acme.org/C=GH/ST=ACC/O=ACME/OU=XS" \  
-reqexts SAN \  
-config <(cat /etc/ssl/openssl.cnf \  
<(printf "\n[SAN]\nsubjectAltName=DNS:acme.org,DNS:*.acme.org")) \  
-out domain.csr
```

⚠ Caution

Make sure that the file generated by the process are only readable by the user running the script.

After you have successfully created a certificate signing request in the file `domain.csr` and a new private key in the file `domain.key`, you need to provide the certificate signing request to the certificate authority of your choice. You must follow the CA-specific process to have the certificate signed and then download the signed certificate as `domain.crt`. Note that you also need all the intermediate certificates and the root certificate which were used by the certificate authority to sign your domain certificate. In the following example, we assume the certificate authority used an intermediate CA certificate "`intermediate.crt`" for signing purposes, and the intermediate certificate was signed by the root CA certificate ("`root.crt`").

The `xs set-certificate` command requires the full certificate chain as one file, and the order of the certificates is important: A signed certificate needs to be followed by the signing certificate. In other words, the domain certificate must be the first certificate, and the root certificate must be the last. You create the certificate-chain file by running the following command:

```
cat domain.crt intermediate.crt root.crt > domain.chain
```

Since the private key must be in PKCS8 format, use the `openssl` command to perform the conversion, as illustrated in the following example:

```
openssl pkcs8 -topk8 -nocrypt -in domain.key -out domain.pk8
```

You can now configure the certificate in XS advanced by running the `xs set-certificate` command, as illustrated in the following example:

```
xs set-certificate example.org -c domain.chain -k domain.pk8
```

→ Tip

The platform performs a connectivity check using the new certificate and returns an error message if it detects that there is something wrong with the certificate. You can remove the files "`domain.chain`", "`domain.pk8`", "`domain.key`", and "`domain.csr`" after completing this step.

To ensure the settings take full effect, restart the XS advanced system by running the `XSA` command as `<sid>adm`, as illustrated in the following example:

```
XSA restart
```

Setting Domain Certificates in PSE Format

It is also possible to use the `set-certificates` command to configure CA-signed domain certificates in the PSE-format, as shown in the following example. The first step is to create a new PSE container together with a certificate signing request. To create the PSE container, use the `sapgenpse` command, as illustrated in the following example:

i Note

The subject and subject alternative names parameters need to be adapted to suit your use case. You should also remember to protect the PSE container with a suitable PIN and ensure that the generated files are only readable by trusted users who really need access to the sensitive data stored in the file.

```
sapgenpse gen_pse -p domain.pse -r domain.csr -k "GN-dNSName:example.org" -k "GN-dNSName:*.example.org" "CN=example.org, C=DE, ST=BW, O=SAP, OU=XS"
```

If the command completes successfully, the certificate signing request will be available in the file "domain.csr", and a new PSE container with the name "domain.pse" is created. Next, you need to provide the certificate signing request to the certificate authority of your choice. You must follow the CA-specific process to have the certificate signed. Request the certificate from the CA as a PKCS7 container; this container includes the signed certificate, all intermediate certificates, and the root certificate.

The next step is to import the signed certificate and its certificate chain into the PSE container. Put the PKCS7 response you receive from your CA in a file called "domain.p7" and import it into the PSE container using the following command:

```
sapgenpse import_own_cert -p domain.pse -c domain.p7
```

You can now configure the certificate using the following command:

```
xs set-certificate example.org --pse domain.pse
```

→ Tip

The platform performs a connectivity check using the new certificate and returns an error message if it detects that there is something wrong with the certificate. You can remove the files "domain.pse", "domain.p7", and "domain.csr" after completing this step.

To ensure the settings take full effect, restart the XS advanced system by running the `XSA` command as `<sid>adm`, as illustrated in the following example:

```
XSA restart
```

Related Information

[Maintaining Trust Certificates in XS Advanced \[page 1706\]](#)

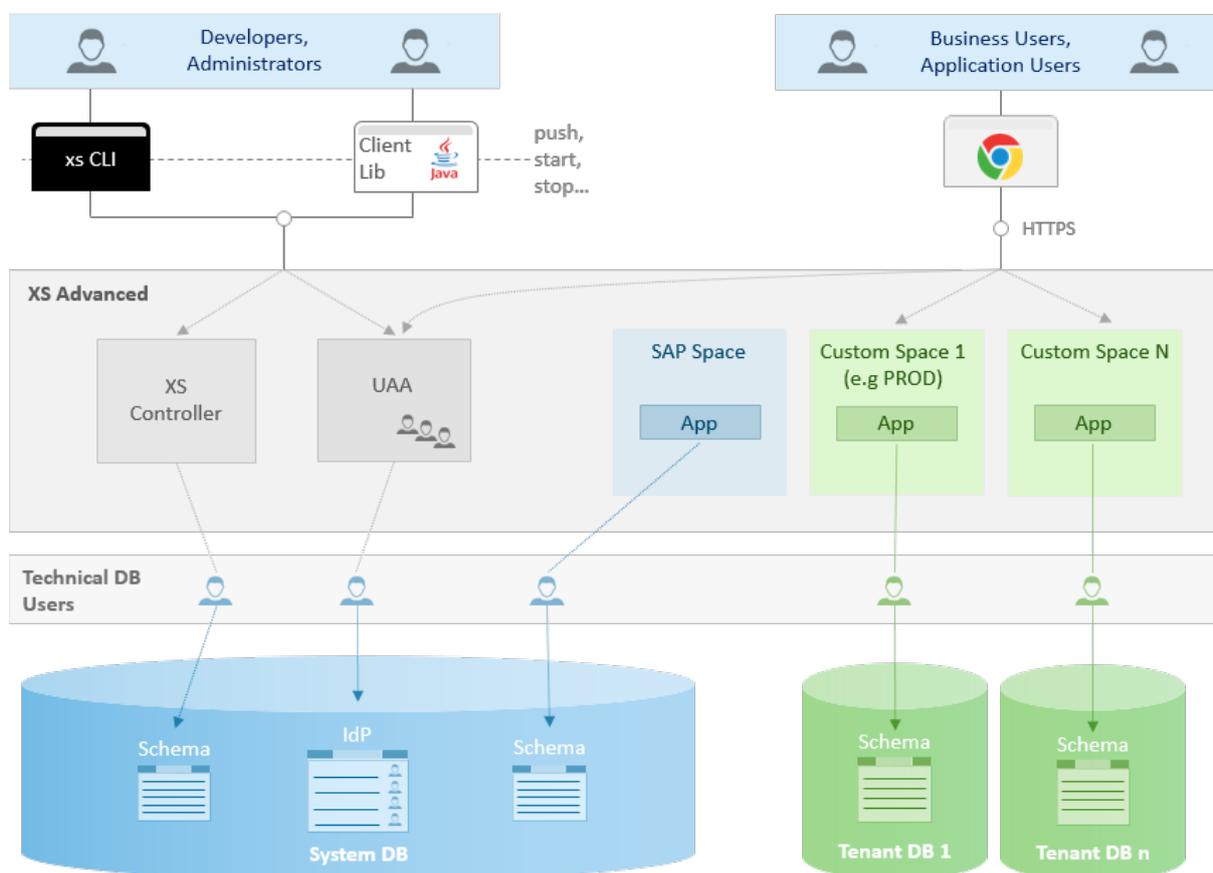
[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

[The XS Command-Line Interface \[page 1653\]](#)

12.2.2.15 Maintaining Tenant Databases in XS Advanced

Use the `xs` command-line interface to maintain tenant database for the XS advanced run-time platform.

This section describes how XS advanced integrates into an multiple-database containers (MDC) setup in SAP HANA. While XS advanced is a system-wide service and uses the System database to persist metadata about applications, XS advanced applications can use an arbitrary tenant database for persistence data. As a result, it is possible to map spaces to one or more tenant databases. One obvious advantage of such a mapping is that service instances created in a mapped space use the tenant database mapped to the space as resource.



XS advanced in an MDC System (Fresh Installation)

- XS Advanced platform services and the User Account and Authentication (UAA) service use the System DB as persistence
- UAA uses the System DB as an Identity Provider (IdP). This means that database users created in the System DB can be used as XS advanced platform users or business users.
- The SAP space is mapped to the System DB by default. This means that XS advanced core applications use the System DB as persistence.
- A custom space (for example 'PROD') is mapped to the default tenant DB. This means that applications bound to SAP HANA service instances in the custom space use the default tenant DB as persistence.

- Newly created spaces are mapped by default to the default tenant DB.

Although there is already a data-isolation concept based on HDI Containers within a single database, using different tenant databases for different spaces enables even stronger isolation of application-specific data.

Viewing Tenant DBs in XS Advanced

As an XS advanced administrator, you can use the `xs tenant-databases` command to display a list of the current mappings between spaces and tenant databases, as illustrated in the following example:

```

Output Code
Viewing Mapped Tenant Database

$ xs tenant-databases

Getting tenant databases as XSA_ADMIN...

name      XSA  status  host           sql-port  mappings
-----
MYTENANT1 no   active  h1.acme.org    30041
MYTENANT2 no   active  h2.acme.org    30044
XSA       yes  active  h3.acme.org    30015    <default>, orgname:PROD
SYSTEMDB  yes  active  h3.acme.org    30013    orgname:SAP

```

The information displayed in the example above includes the following details:

XS Advanced Application Details

Property	Description
name	The name of the database tenant
XSA	The listed tenant database is enabled (“yes”) or not yet enabled (“no”) for use with XS advanced
status	An <i>active</i> tenant database is up, running, and available; a tenant database that is “inactive” is either stopped, in the process of being created, or otherwise unavailable.
host	The name of the master host on where the tenant database is running
sql-port	The port to use for SQL connections to this tenant database's master host
mappings	A comma-separated list of the XS advanced spaces (<code>org:space</code>) that are mapped to the tenant database

In the example describes the following mapping configuration:

- The tenant databases `MYTENANT1` and `MYTENANT2` are active, but not yet enabled for use with XS advanced.
- The tenant database `XSA` is enabled for use in XS advanced and the tenant is mapped by default when a new XS advanced space is created. In addition, the tenant database `XSA` is mapped explicitly to the space `PROD` in the organization `orgname`.

- The `systemDB` is used by applications in space `SAP` within the organization `orgname`.

Enabling Tenant DBs for XS Advanced

To register a new tenant database, for example, `MYTENANT1`, for use in XS advanced use the following command:

```
$ xs enable-tenant-database MYTENANT1
```

If you use the `xs` CLI to create a new tenant database, the new database is immediately **enabled** for use with XS advanced, as shown in the following example:

Output Code

```
$ xs tenant-databases
```

```
Getting tenant databases as XSA_ADMIN...
```

name	XSA	status	host	sql-port	mappings
MYTENANT1	yes	active	h1.acme.org	30041	
MYTENANT2	no	active	h2.acme.org	30044	
XSA	yes	active	h3.acme.org	30015	<default>, orgname:PROD
SYSTEMDB	yes	active	h4.acme.org	30013	orgname:SAP

→ Tip

The `XSA` column indicates if a tenant database is either enabled (`yes`) or disabled (`no`) for use in XS advanced. A **disabled** database also does not have either a host or an SQL port assigned.

Alternatively, you can use the `xs` command `xs create-tenant-database`, which creates a tenant DB and registers it for use with XS advanced automatically:

```
$ xs create-tenant-database MYTENANT3
```

Mapping Tenant DBs in XS Advanced

After a tenant database is enabled for use with XS advanced, you can map organizations and spaces to the newly enabled tenant database with the `xs map-tenant-database` command, as shown in the following example:

```
$ xs map-tenant-database -o orgname -s myspace MYTENANT1
$ xs map-tenant-database -o myorg2 MYTENANT2
```

→ Tip

If you map a tenant database to an organization, the mapping is inherited by **all** spaces in the mapped organization, as long as no tenant is set explicitly.

To display an overview of the mappings between tenant databases and organizations and spaces, use the `xs tenant-database-mappings` command, as shown in the following example:

Output Code

```
$ xs tenant-database-mappings
Getting tenant database mappings as XSA_ADMIN...

organizations and spaces      tenant databases
-----
myorg2                        MYTENANT2
myorg2:dev                    MYTENANT2 (inherited from orgname)
orgname                       XSA (inherited from global)
orgname:PROD                  XSA
orgname:SAP                   SYSTEMDB
orgname:myspace               MYTENANT1
```

If the same space or organization is mapped to multiple tenant databases, you can use the `--default` option to specify one of the tenant databases as the default database, as shown in the following command:

```
$ xs map-tenant-database -o orgname -s myspace MYTENANT2 --default
```

The selected **default** tenant database is flagged in the output of the `xs tenant-database-mappings` command as shown in the following example:

Output Code

```
$ xs tenant-database-mappings
Getting tenant database mappings as XSA_ADMIN...

organizations and spaces      tenant databases
-----
[...]
orgname:myspace               MYTENANT1, MYTENANT2 (default)
[...]
```

If no default tenant database is declared or you want to use a particular mapped tenant database for a certain SAP HANA service instance, you must specify the GUID of the target database whenever you create an SAP HANA service in the target space, as illustrated in the following example:

```
$ xs create-service hana hdi-shared myservice -c '{"database_id":
"0e140209-5e3a-42b3-8c7a-21260eba4287"}'
```

To display a list of know tenant databases and the corresponding database GUIDs use the `--guids` option, as shown in the following example:

Output Code

```
$ xs tenant-databases --guids
Getting tenant databases as XSA_ADMIN...

name      XSA  status  host           sql-port  id              mappings
-----
MYTENANT1 yes   active  h1.acme.com    30041    0e140209-...   orgname:myspace
```

Unmapping and Deleting Tenant DBs in XS Advanced

To unmap a tenant database from a space or organization, use the command `xs unmap-tenant-database`, as shown in the following example:

```
$ xs unmap-tenant-database -o orgname -s myspace MYTENANT1
```

To disable a tenant DB for use with XS advanced, use the command `xs disable-tenant-database`, as shown in the following example:

```
$ xs disable-tenant-database MYTENANT1
```

To delete a tenant database, use the command `xs delete-tenant-database`. You need to specify the credentials for the system users of both the System database (`SYSTEMDB`) and the specified tenant database (`MYTENANT2`), as shown in the following example:

```
$ xs delete-tenant-database MYTENANT2 -n <TENANTDB_SYSTEM_USER> -t <PASSWD> -u  
<SYSTEMDB_SYSTEM_USER> -p <PASSWD>
```

⚠ Caution

Deleting a tenant database destroys all data within the specified tenant database.

Related Information

[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

[The XS Command-Line Interface \[page 1653\]](#)

12.2.2.16 Maintaining Host Pinning

Use the `xs` command-line interface to maintain services for applications running in the XS advanced run-time platform.

Host pinning enables you to “pin” an XS advanced application to a particular set of hosts that have been assigned the role `xs_worker`. In this way, you can control where applications run and ensure they can only be run on a particular host or in a specific space. Pinning applications to a particular host provides the following benefits:

- Isolation
Ensure that applications run inside the same SAP HANA system but completely isolated from one another on specific (different) hosts.
- Sizing
Exploit application-specific hardware and sizing benefits which are only available on particular hosts

XS advanced distinguishes between the following host pinning modes:

- `strict`

The pinned application will only start if one of the pinned hosts is available. (Default)

- relaxed

The pinned application will start even if none of the pinned hosts is available.

Viewing Currently Pinned Hosts

You can display a list of the applications that are currently pinned in the target space. The following example shows that no applications are pinned to any host or space, and no **default** pinning configuration has been set up yet:

Output Code

```
$ xs pinned-hosts

Getting apps in org "orgname" / space "SAP" as XSA_ADMIN...

Found apps:

app                pinned to      mode
-----
auditlog-db        <not set>
auditlog-server    <not set>
auditlog-broker    <not set>
deploy-service     <not set>
component-registry-db <not set>
product-installer  <not set>
auditlog-odata     <not set>
auditlog-ui        <not set>

Default host pinning for target space "SAP": <no pinning>
```

Viewing Hosts Available for Pinning

To pin applications to certain hosts, you first need to find out which hosts are available for pinning duties, for example, by using the command `xs system-info`, which displays a list of all hosts running an execution agent (assigned the `xs_worker` role) along with some additional details, as illustrated in the following example:

Output Code

```
$ xs system-info

Getting system infrastructure information...

Execution agents:
-----
1. host1.acme.org:47588
   created at      Dec 15, 2017 7:50:48 AM
   port range      50000-50499
   os.arch         amd64
   os.name         Linux
   java.vendor     SAP AG
   java.version    1.8.0_144
   os.version      3.12.74-60.64.40-default
```

```

sun.arch.data.model 64
version             v1.0.70

2. host2.acme.org:63506
  created at        Dec 15, 2017 8:25:08 AM
  port range        50000-50499
  os.arch           amd64
  os.name           Linux
  java.vendor       SAP AG
  java.version      1.8.0_144
  os.version        3.12.74-60.64.40-default
  sun.arch.data.model 64
  version           v1.0.70
[...]
```

The system output shown in the example above shows two running Execution Agents (index: 1 and 2), which can be used to pin applications to the hosts where the agents are registered.

Pinning Applications to Hosts

It is possible to pin either specific applications or entire spaces to a particular host or hosts. To pin an application to specific hosts, use the command `xs pin-hosts <application> <host>`. There are two ways to specify the respective host when using the `xs pin-hosts` command:

- Specify the Execution Agent index
The index is the number associated with the hosts listed by the `xs system-info` command - "1" or "2" in the example above.
- Specify the host name
The host name is listed along side the "index" in the output displayed by the `xs system-info` command.

The following two commands are equivalent and pin the application `auditlog-ui` to the host `host2.acme.org`:

```
$ xs pin-hosts auditlog-ui --eas 2
```

```
$ xs pin-hosts auditlog-ui --hosts host2.acme.org
```

Assuming you are logged on to the same XS advanced system where you ran the `xs system-info` command in the example, above, the following command pins the `deploy-service` application to the same host where Execution Agent 2 is running, but in `relaxed` pinning mode:

```
$ xs pin-hosts deploy-service --eas 2 --relaxed
```

→ Tip

To change an existing host-pinning setting, use the `--unset-pinning` option.

To pin a space to a host, you can use the `xs update-space` command, which works in a similar way to the `xs pin-hosts` command. For example, to pin the space "SAP" to the host `host2.acme.org`, run the following command:

```
$ xs update-space SAP --hosts host2.acme.org
```

To display details of all currently pinned hosts, use the `xs pinned-hosts` command, as illustrated in the following example:

Output Code

```
$ xs pinned-hosts

Getting apps in org "orgname" / space "SAP" as XSA_ADMIN...
Found apps:

app                pinned to          mode
-----
auditlog-db        <not set>
auditlog-server    <not set>
auditlog-broker    <not set>
deploy-service     host2.acme.org    relaxed
component-registry-db <not set>
product-installer  <not set>
auditlog-odata     <not set>
auditlog-ui        host2.acme.org    strict

Default host pinning for target space "SAP":
[host2.acme.org] (strict)
```

Note

To apply and enable the new (or modified) host-pinning configuration, you must restart the pinned applications. Explicit host-pinning settings for an application takes precedence over the host-pinning settings for the space in which the application is deployed and running.

Troubleshooting Host-Pinning Problems

Stopping the Execution Agent on a host (for example, `host2.example.org`) by removing the `xs_worker` role from a host illustrates the difference between “strict” and “relaxed” pinning. A short time after the Execution Agent stops, the state of the pinned applications look like in the following example:

Output Code

```
$ xs apps

Getting apps in org "acme" / space "SAP" as XSA_ADMIN...
Found apps:

name                requested  instances  memory  urls
                   state
-----
[...]
deploy-service     STARTED   1/1        280 MB  https://deploy-service.ex[...]
[...]
auditlog-ui        STARTED   0/1        64.0 MB  https://auditlog-ui.examp[...]
```

The output from the `xs apps` command shows in the `instances` column that the `auditlog-ui` application could not start; this is because the pinning mode was set to “strict”. The `deploy-service` command,

however, was able to start because the pinning mode was “relaxed”. The logs for the `auditlog-ui` application contains the information illustrated in the following example:

Output Code

```
$ xs logs auditlog-ui --last 2

-----
Connected, dumping recent logs for app "auditlog-ui"
12/15/17 [...PM] [API] ERR  Number of running instances for app 'auditlog-ui'
12/15/17 [...PM] [API] ERR  Failed to start an instance of app 'auditlog-ui'
```

Related Information

[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

[The XS Command-Line Interface \[page 1653\]](#)

[The XSA Command Reference \[page 1740\]](#)

12.2.2.17 Maintaining Build Packs in XS Advanced

Build packs transform application code into an application droplet that can be deployed to the appropriate run-time environment.

A build pack is the component that XS advanced uses to transform application code into an application droplet, which can be run as an application instance on the XS advanced platform. This transformation process is called “staging”. When pushing an application to the run-time environment, the appropriate build pack is automatically detected for the pushed application, and the application is staged using the selected build pack. For example, if you push a WAR file, the Java build pack detects that it is the appropriate build pack to create a usable Tomcat application from that WAR file. Usually a build pack adds the corresponding run time environment (for example, Node.js or Tomcat) along with any specified dependencies (for example, Node modules) to the application being pushed, and it may even add additional libraries (for examples, JDBC drivers) and configure the application to connect to bound services.

Viewing Installed Build Packs

By default XS advanced comes with a set of default build packs that are supported by SAP. The command `xs buildpacks` shows a list of all the build packs installed on the XS advanced platform. The following example shows a list of all default build packs:

```
$ xs buildpacks
Getting buildpacks...
buildpacks          version  position  enabled  locked
-----
sap_java_buildpack  1.6.23   1         true    false
sap_nodejs_buildpack 3.4.3    2         true    true
```

```
sap_python_buildpack 0.2.2 4 true
```

The names of the build packs are generally self-explanatory; in the example above, the build packs names indicate support for Java, Node.js, and Python applications.

The value for the position attribute displayed in the output of the `xs buildpacks` command is important if the build pack needs to be detected automatically. This is because the detection procedures for the various build packs are executed in the order of their positions; in this example: position 1, 2, or 4. The first build pack to report that it has detected an application it thinks it is able to stage is used. Build packs that are `enabled` can be used for staging an application; disabled build packs are not able to detect any uploaded files or indicate that they are responsible for staging a particular type of application. The default setting for `enabled` is `"true"`. Locking a build pack ensures that the locked build pack stays at the current version; it cannot be updated with a new version. The default setting for `locked` is `"false"`.

Creating a Build Pack

You can use the `xs create-buildpack` command to create your own custom build pack. Creating your own build pack is useful if you want to run applications in XS advanced, which are written in a programming language that is not supported by the default build packs. For example, if you want to run a Go application in XS advanced, you will need to write your own custom build pack for Go. For more information about creating custom build packs, see the section *"Custom Build Packs in XS Advanced"* in the *SAP HANA Developer Guide for XS Advanced* (in *Related Information* below). After you have implemented your custom build pack, you can upload it to the XS advanced platform by running the following command, which would place it at position 5 in the build pack detection process:

```
xs create-buildpack my_custom_buildpack /path/to/my_custom_buildpack 5
```

Updating a Build Pack

You can update a build pack by using the `"xs update-buildpack"` command, for example:

```
xs update-buildpack my_custom_buildpack -p /path/to/version2/my_custom_buildpack
```

The default build packs provided by SAP are automatically updated with new features and security patches, when you update the XS advanced platform.

Updating a build pack does not directly affect any applications that were staged with the old version of the build pack. If a security patch or a new feature of the build pack is applied to an application, the application needs to be restaged. Restaging happens when the application is deployed or pushed. A restage, however, can also be triggered individually. After restaging an application, it is necessary to restart the application, too, in order to ensure the execution of the new droplet produced by the new version of the build pack, as illustrated in the following example:

```
xs restage <app_name>
xs restart <app_name>
```

It is also possible to enforce a restage of all applications. However the procedure requires a restart of XS advanced. To trigger a restart of all application, log on to XS Advanced as the `<sid>adm` user and run the following (XSA) commands:

```
XSA restage-at-startup
XSA restart
```

Using Git Build Packs

Creating and updating a build pack as described in this section, requires administrator privileges within the XS advanced platform. However, when pushing an application, it is also possible to specify a git repository that contains the source code of a build pack, as illustrated in the following example

```
xs push my-application -f my-manifest.yml -b https://my-git-server.example.com/
my_custom_buildpack.git
```

In this scenario, the build pack is downloaded from the specified git repository and used when staging the pushed application. Note that it is only possible to specify a Git repository URI that is also valid for the `git clone` command. The URI can use either the `git` or the `https` protocol and needs to be reachable from the XS Controller host, and the source code of the build pack must to be available in the default branch of the repository. Note that if you are calling a Git server via HTTPs, it might be necessary to configure the correct trust certificates first. For more information about how to configure trust certificates within XS advanced, see *Maintaining Trust Certificates* in *Related Information* below.

Related Information

[Maintaining Trust Certificates in XS Advanced \[page 1706\]](#)

12.2.2.18 Maintaining Service URLs

The SAP HANA XS advanced platform supports the creation and maintenance of a list of service URLs, which consist of a service name and the corresponding URL. Some service URL mappings are created during installation of XS advanced, but additional URLs can be registered by an administrator at any point in time. Usually, the service URL points to an application running on the XS advanced platform, but the URL could also point to an external service.

This list of registered service URLs can be read by client tools, for example, the `xs` command-line tool without requiring any logon credentials. If the XS Controller API URL is set, the known service URLs are displayed in the output of the `xs version` command, as illustrated in the following example:

☰ Output Code

Registered Service URLs

```
$ xs version
...
-----
Registered service URLs:
  deploy-service           = https://deploy-service.xsa.example.com:35033
  product-installer       = https://orgname-sap-product-installer.xsa.examp...
  job-scheduler-dashboard = https://jobscheduler-dashboard.xsa.example.com:350
  product-installer-ui    = https://orgname-sap-product-installer-ui.xsa.ex...
  xsa-cockpit             = https://xsa-cockpit.xsa.example.com:35033
```

Links to the services are also created on the status page that is delivered to a Web browser that is pointed to the XS advanced API URL. In addition, the platform provides a redirection mechanism for the registered services. Pointing a Web browser to `<API-URL>/go/xsa-cockpit` redirects the browser to the URL registered for the service `xsa-cockpit`. This is particularly useful if you want to provide stable and predictable links to front-end applications even in **port**-routing mode without having to create routes with static ports.

Registering Service URLs

Service URLs for applications deployed on the XS advanced platform can be registered automatically using the corresponding feature in the application's `mtad.yaml` deployment descriptor. XS advanced administrators can, however, also register a service URL manually, using the `xs register-service-url` command, as illustrated in the following example:

☰ Output Code

Register a Service URL

```
$ xs register-service-url myservice https://myservice.xsa.example.com
```

To display a list of all currently registered service URLs, run the following command:

☰ Output Code

List all Registered Service URLs

```
$ xs service-urls
...
-----
Registered service URLs:
  deploy-service           = https://deploy-service.xsa.example.com:35033
  product-installer       = https://orgname-sap-product-installer.xsa.exempl...
  job-scheduler-dashboard = https://jobscheduler-dashboard.xsa.example.com:3503
  product-installer-ui    = https://orgname-sap-product-installer-ui.xsa.exa...
  xsa-cockpit             = https://xsa-cockpit.xsa.example.com:35033
```

Unregistering Service URLs

Service URLs can be manually removed from this list of registered service URLs by using the `xs unregister-service-url` command, as illustrated in the following example:

Output Code

Unregister a Service URL

```
$ xs unregister-service-url myservice
...
Removed registration of service name "myservice" with URL "https://
myservice.xsa.example.com".
OK
```

Related Information

[Maintaining XS Advanced Run-Time Components with the XS CLI \[page 1653\]](#)

12.2.2.19 Building Roles for XS Advanced Applications

Use authorization artifacts to control access to XS advanced applications.

In XS advanced, application developers create and deploy application-based authorization artifacts for business users. Administrators use this information to build roles, define sets of roles called role collections, and assign these collections to business users or user groups. In this way, they control the users' permissions and, as a result, access to the applications.

Role Templates

Applications define role templates acting as a blueprint for real-life role instances. Based on these role templates, administrators create role instances by filling in concrete values for attributes defined by role templates. For example a role template for editing actions on HR data could have an attribute to configure the region or country to restrict the editing permission to. To display the role templates provided by applications in the current target space or made available globally, use the command `xs role-templates`, as illustrated in the following example:

Output Code

Role Templates for XS Advanced Applications

```
$ xs role-templates

Getting role templates in space "SAP" as user "XSA_ADMIN"...

app                role template    attributes        description
```

```

-----
alm                ControllerAdmin                Template for Ctrl Admin
xs_role_admin     XS_ROLE_ADMIN                Auths for XS role build
fileprocessor     Admin                        Role for Admin UI
                  API                          Role for API access
                  Auditor                       Role for read-only acce
auditlog-ui       AuditLogViewer                View all auditlogs
java-hello-world  Viewer                        Country                View all books
                  Editor                          Country, CostCenter    Edit, delete the books

```

In the example above, the role templates “Viewer” and “Editor” of the application `java-hello-world` are examples for role templates with attributes (`Country` and `CostCenter`). It is not possible to assign role templates directly to users; a specific instance of the role template must be created first, in the form of a role, which is described in more detail in the next section.

Roles in XS Advanced Applications

If you are basing a role on a role template that does not specify any attributes, creating the role is a straightforward process, for example, using the `xs create-role` command, as illustrated in the following example:

```

$ xs create-role fileprocessor Admin FP-Admin "Role for fileprocessor Admin UI
access"
Creating role "FP-Admin" for app "fileprocessor" using role template "Admin" as
user "XSA_ADMIN"...
OK

```

The example command above creates a new role called `FP-Admin` from the `Admin` template of the `fileprocessor` application. No attributes are specified in this example, but it is always a good idea to provide a description for the new role; this information is used when the role instance is created and appear in the list of roles that can be displayed with the `xs roles` command, as illustrated in the following example:

Output Code

Displaying XS Advanced Application Roles

```

$ xs roles

Getting application roles in space "SAP" as user "XSA_ADMIN"...
Found roles:

name                app                template            descript
-----
XS_CONTROLLER_ADMIN <cloud_controller> XS_CONTROLLER_ADMIN Default
XS_CONTROLLER_AUDITOR <cloud_controller> XS_CONTROLLER_AUDITOR Default
XS_CONTROLLER_USER   <cloud_controller> XS_CONTROLLER_USER   Default
XS_AUTHORIZATION_ADMIN <xs_authorization> XS_AUTHORIZATION_ADMIN Default
XS_AUTHORIZATION_DISPLAY <xs_authorization> XS_AUTHORIZATION_DISPLAY Default
XS_MONITOR_ADMIN     <xs_monitor>       XS_MONITOR_ADMIN     Default
XS_MONITOR_DISPLAY   <xs_monitor>       XS_MONITOR_DISPLAY   Default
XS_SUBSCRIPTION_ADMIN <xs_subscription> XS_SUBSCRIPTION_ADMIN Default
XS_SUBSCRIPTION_DISPLAY <xs_subscription> XS_SUBSCRIPTION_DISPLAY Default
XS_TENANT_ADMIN      <xs_tenant>       XS_TENANT_ADMIN      Default
XS_TENANT_DISPLAY    <xs_tenant>       XS_TENANT_DISPLAY    Default
XS_USER_ADMIN        <xs_user>         XS_USER_ADMIN        Default
XS_USER_DISPLAY      <xs_user>         XS_USER_DISPLAY      Default
XS_USER_PUBLIC       <xs_user>         XS_USER_PUBLIC       Default
ControllerAdmin     alm                ControllerAdmin      Default
XS_ROLE_ADMIN       xs_role_admin     XS_ROLE_ADMIN        Default

```

AuditLogViewer	auditlog-ui	AuditLogViewer	Default
API	fileprocessor	API	Default
Admin	fileprocessor	Admin	Default
Auditor	fileprocessor	Auditor	Default
FP-Admin	fileprocessor	Admin	Admin UI

The `xs roles` command shows all application roles scoped for the current target space and any global roles. In addition to roles created from application-defined role templates, there are also some roles defined by the platform itself. Platform-defined roles typically have names starting with "xs_" and the origin of the role is shown in angle brackets <...> to distinguish them from normal application roles.

When creating a role from a role template with attributes, the `xs create-role` command prompts the user to specify the values to use for all the attributes. Alternatively, the attribute values can be provided on the command line with the option `-a ATTRIBUTES` where `ATTRIBUTES` can either be a JSON format structure with all attribute values, or the path to a text file that contains the JSON structure. A convenient way to create such a file is by letting the `xs` client create it for a given role template, as illustrated in the following example:

Output Code

Creating a Sample `attributes.json` File for a Role Template

```
$ xs create-role java-hello-world Editor --create-sample attributes.json
Creating attribute values sample file for role template "Editor" for app
"java-hello-world"...
OK
```

The resulting file `attributes.json` can then be edited in any text editor to fill in the desired values. To create a new role with these attributes and the edited values, use the option `-a attributes.json` as illustrated in the following example:

Output Code

Creating an Application Role Using Attribute Values Stored in a JSON File

```
$ xs create-role java-hello-world Editor Editor-Germany -a attributes.json
"Editor role for Germany"
Creating role "Editor-Germany" for app "java-hello-world" using role template
"Editor" as user "XSA_ADMIN"...
OK
```

The details of the newly created (or any other) role can be displayed using the `xs role` command. The details include a description of the role, the provided scopes, and any attribute values, as illustrated in the following example:

Output Code

Displaying Details of an Application Role

```
$ xs role Editor-Germany

Getting role "Editor-Germany" in space "SAP" as user "XSA_ADMIN"...

name:           Editor-Germany
description:    Editor role for Germany
app:            java-hello-world
template:       Editor (Edit and Delete the books)

scope           description
```

```

-----
java-hello-world!il.Create      create
java-hello-world!il.Delete     delete

attribute  value      source  description
-----
Country    Germany  static  Country
CostCenter de-01     static  CostCenter

```

Role Collections in XS Advanced

In XS advanced, roles are not directly assigned to users; they are managed in reusable role collections that can, for example, list all the roles required to perform a specific task. New role collections can be created using the `xs create-role-collection` command, as illustrated in the following example:

Output Code

Creating a Role Collection

```

$ xs create-role-collection Example-Collection "An example role collection"
Creating role collection "Example-Collection" as user "XSA_ADMIN"...
OK

```

To display existing role collections, use the command `xs role-collections`, as illustrated in the following example:

Output Code

Displaying a List of Role Collections

```

$ xs role-collections

Getting role collections as user "XSA_ADMIN"...

role collection      description
-----
AUDITLOG_VIEWER
XS_AUTHORIZATION_ADMIN      Authorizations for XS role builder
XS_AUTHORIZATION_DISPLAY    Authorizations for XS role viewer
XS_USER_ADMIN               Admin authorizations for XS user management
XS_USER_DISPLAY             Display authorizations for XS user management
XS_USER_PUBLIC              Default authorizations for XS user
XS_MONITOR_ADMIN            Authorizations for XS monitoring management
XS_MONITOR_DISPLAY          Authorizations for XS monitoring display
XS_SUBSCRIPTION_ADMIN       Authorizations for XS subscriptions management
XS_SUBSCRIPTION_DISPLAY     Authorizations for XS subscriptions display
XS_TENANT_ADMIN             Authorizations for XS tenants management
XS_TENANT_DISPLAY           Authorizations for XS tenants display
XS_CONTROLLER_ADMIN         Authorizations for XS controller admin
XS_CONTROLLER_USER          Authorizations for XS controller user
XS_CONTROLLER_AUDITOR       Authorizations for XS controller auditor
Example-Collection          An example role collection

```

You can add individual roles to (or remove them from) a role collection using the command `xs update-role-collection`, as illustrated in the following example:

Output Code

Adding a Role to a Role Collection

```
$ xs update-role-collection Example-Collection --add-role FP-Admin
Getting role "FP-Admin" in space "SAP" as user "XSA_ADMIN"...
Updating role collection "Example-Collection" as user "XSA_ADMIN"...
OK
```

To show the contents of a role collection, use the command `xs role-collection`, as shown in the following example:

Output Code

Displaying Details of a Role Collection

```
$ xs role-collection Example-Collection
Getting role collection "Example-Collection" as user "XSA_ADMIN"...
Roles of role collection "Example-Collection" (An example role collection):
name                app                space                template            description
-----
FP-Admin            fileprocessor      orgname:SAP         Admin              Role Admin UI access
Editor-Germany      java-hello-world  orgname:SAP         Editor             Edit role for Germany
```

Assigning Role Collections

The final step in role management is assigning the role collection to users. The XS command-line client provides the `assign-role-collection` command for this purpose, as shown in the following example:

Output Code

Assigning a Role Collection to an XS Advanced User

```
$ xs assign-role-collection Example-Collection MYUSER
Assigning role collection "Example-Collection" to user "MYUSER"...
OK
```

Use the `xs unassign-role-collection` command to remove any role collections assigned to users. To display the role collections currently assigned to a user, use the `xs assigned-role-collections` command, as illustrated in the following example:

Output Code

Displaying Details of Assigned Role Collections

```
$ xs assigned-role-collections MYUSER
Getting role collections assigned to user "MYUSER"...
```

role collection	description
XS_USER_PUBLIC	Default authorizations for XS user
Example-Collection	An example role collection

Related Information

[Maintaining Platform Users in XS Advanced \[page 1671\]](#)

12.2.2.1.20 Maintaining XS Advanced Application Run Times

Keep XS advanced application run-time versions up to date.

Applications run times are used by buildpacks to compile an application during the staging process. XS advanced keeps a store for application run times such as virtual machines for Java, Node.js, application containers, and so on. The application run-time store is automatically updated with most recent version of the application run time when XS advanced is updated to a new version. However, the XS advanced administrator can also maintain application run times manually.

Viewing Installed Application Run Times

To display a list of the installed run-time environments in XS advanced, use the `xs runtimes` command, as illustrated in the following example:

Output Code

Listing Installed Run-Time Environments

```
$ xs runtimes
```

```
Getting runtimes...
```

type	version	id	resolved	active	description	bound apps
hanajdbc1	120.20	7	true	true	SAP HANA JDBC Driver 1.120.20	0
hanajdbc2	3.19	6	true	true	SAP HANA JDBC Driver 2.3.19	4
node6.12	2.1	3	true	true	Node.js 6.12.2.1 for Linux x86...	0
node8.9	3.5	2	true	true	Node.js 8.9.3.5 for Linux x86-...	10
sapjvm8	1.35	4	true	true	SAP JVM 8 Patchlevel 35 for Liux	0
sapjvm8_jre	1.35	5	true	true	SAP JVM JRE 8 Patchlevel 35 fo...	4
tomcat8	5.23	0	true	true	Apache Tomcat Web Container 8....	4
tomee1.7_jaxrs	5	1	true	true	Apache TomEE jaxrs 1.7.5	0

For each application run time environment, the following information is displayed:

Application Run-Time Environment Details

Information	Description
type	The name of the application run-time environment
version	The version of the application run-time environment
id	The index of the application run time by which it is referenced in other xs CLI commands (for example, <code>xs runtime</code>)
resolved	The indicated application run time is correctly detected by the XS advanced run time environment [true false]
active	The application run time is enabled for use by build packs [true false]
description	A short summary of the application run-time details
bound apps	The number of applications currently using the application run time

Finding Bound Applications

To display a list of all applications that use a particular application run time environment, use the `xs runtime` command and specify a particular run time with the run-time index, as illustrated in the following example:

→ Tip

The run-time index is displayed in the `id` column of the output of the `xs runtimes` command.

≡ Output Code

Listing Applications Bound to a Specific Run Time

```
$ xs runtime -i 2

Showing information about "node8.9"

type:          node8.9
version:       3.5
id:            2
description:   Node.js 8.9.3.5 for Linux x86-64
resolved:      true
active:        true
bound apps:    auditlog-db, auditlog-broker, component-registry-db, auditlog-ui,
               jobscheduler-db, jobscheduler-rest, jobscheduler-service,
               jobscheduler-backend, jobscheduler-dashboard, jobscheduler-broker
```

You can use the information displayed by the `xs runtime` command to check which applications need to be restaged if a new application run time has been installed and, in addition, to ensure that all relevant applications are using the specified application run-time environment.

Pinning Run Time Environments to Applications

Normally, the buildpack chooses the latest version of an application run time during application staging. However, in support cases, you can pin a specific version of a run time to an application, forcing the buildpack to use the specified run-time version during application staging. For example using the command `xs pin-runtime` and the options `<ALIAS>`, `<APP>`, and `-i <PINNED RUNTIME>`, you can ensure that the application run time with index `<PINNED RUNTIME>` is always used when an application build pack requests a run time with the name `<ALIAS>`.

The `<ALIAS>` can either be an exact run-time type (for example, "node8.9" to pin a fixed run-time version if the application build pack requires the Node.js version 8.9) or a prefix of a run-time type (for example, "node8" to pin a fixed run-time version if the application build pack requires any version of Node.js version 8.*).

⚠ Caution

Support for run-time pinning is provided for Java and Node.js types only.

In the following example, a Node.js application is pinned to a Node.js "6" run time **type** despite the fact that the buildpack would normally choose a Node.js run time of type "8":

→ Tip

The run-time type is displayed in the `type` column of the output of the `xs runtimes` command.

≡ Output Code

Pinning Applications to a Run-Time Type

```
$ xs pin-runtime node8 myapp -i 3
Updating app "myapp" in org "myorg" / space "PROD" as XSA_ADMIN...
```

To ensure that the specified changes take effect, the target application must be re-staged and re-started. You can use the `xs pinned-runtimes` to display a list of the mappings between run times and applications, and filter the output for a specific application name, as illustrated in the following example:

≡ Output Code

Displaying Pinned Run Times

```
$ xs runtimes myapp
Showing pinned runtimes for app "myapp".
alias  type      version id resolved active description                                bound
-----
node8  node6.12  2.1    3   true    true    Node.js 6.12.2.1 Linux x86...                0
```

To undo the changes, and revert to the previous run-time settings, use the `xs unpin-runtime` command, as illustrated in the following example:

☰ Output Code

Removing Application-to-Run-Time Pinnings

```
$ xs unpin-runtime node8 myapp  
Updating app "myapp" in org "myorg" / space "PROD" as XSA_ADMIN...
```

Creating and Uploading a New Run Time

You can upload a new run time to XS advanced with the command `xs create-runtime`; you need to specify the directory or .zip file containing the application run time that you want to upload, as illustrated in the following example:

```
$ xs create-runtime -p Python-2.7.6
```

! Restriction

Only application run times supported by XS advanced can be uploaded, for example: SAP JVM, SAP Node.js, Python, Tomcat, TomEE, and HANA JDBC driver.

Updating Run Times

When updating XS advanced, the default application run times are automatically maintained and updated as well, for example, to provide security patches.

i Note

Manually uploaded application run times are not included in the automatic update.

To ensure that all applications use the updated application run times, all applications must be restaged and restarted. You can use the commands `xs restage` and `xs restart` to restage and restart individual applications manually, or use the `XSA` command to restage **all** applications during XS advanced startup, as illustrated in the following example:

i Note

Log on to the SAP HANA system as `<sid>adm`.

```
$ XSA restage-at-startup  
$ XSA restart
```

⚠ Caution

These XSA commands restart not only XS advanced but also all applications.

Related Information

[Maintaining Applications in XS Advanced \[page 1684\]](#)

12.2.2.1.21 Maintaining Tasks in XS Advanced

Schedule and manage one-off tasks with XS advanced applications.

XS advanced applications are Web applications that on startup open a port and then remain available to respond to user requests until the applications are stopped. This type of application life-cycle is not appropriate for all types of work. For example, it is not appropriate for initialization or setup tasks that need to be performed only once and do not require any user interaction; these types of problems are best handled by using so called one-off tasks. One-off tasks are special application instances that are started by specifying a particular custom command on application startup. No route is created for a one-off task, which means that the custom command can perform any type of activity that does not require user interaction. After the custom command exits, the one-off tasks finishes and reports the command's exit code to the user.

Running a One-Off Task

Before using an application to perform a one-off task, the corresponding application must first be deployed, as described in *Maintaining Applications in Related Information*. After the application has been successfully deployed, the command `xs run-task` can be used to create a new task instance to run it. The following arguments have to be used to specify the required information in the `run-task` command:

1. The name of the application (and therefore the droplet files for the task execution)
2. A name for the new task
3. The command to execute

The following example shows what a typical task start command looks like:

≡ Output Code

```
$ xs run-task my-app my-task "echo Hello World" --wait-for-completion
Running task "my-task" on app "my-app"...
Task finished successfully.
```

The new task "my-task" is started based on the current droplet and environment variable settings of the application "my-app". Any additional environment variable values that are required can be provided by using the option `-e` with the command `run-task`. The option `--wait-for-completion` forces the client to wait for the task execution to finish. With this option, the success (or failure) of the task execution is reported at the end; without this option, the task execution is started asynchronously and the `run-task` command returns immediately.

→ Tip

You can display the current state of a task by using the `xs tasks` command.

Checking the Current State of Tasks

The command `xs tasks` is used to display a list of existing tasks for a given application and check their current state, as shown in the following example:

Output Code

```
$ xs tasks my-app
Listing tasks for app "my-app"...
-----
index  name      created                command                state
-----
7      bad-task  Apr 25, 2018 10:07:43 AM non-existing command  FAILED
      Failure reason: Process terminated with exit code 1
8      my-task   Apr 25, 2018 10:08:22 AM echo Hello World      SUCCEEDED
```

Although failed tasks have an additional short failure reason message, more useful information can be found in the task-execution output.

Displaying Task Output

The task execution is very similar to the execution of regular apps. Therefore, the output of the task can be found in the logs of the corresponding application. As documented in [Displaying Application Logs](#), the command `xs logs` is used to display the application logs. For the example above, the log output could look like this:

Output Code

```
$ xs logs my-app --recent
Connected, dumping recent logs for app "my-app"...
-----
4/25/18 10:07:43.385 AM [API] OUT Starting new task instance '39660f9d-6eed...
4/25/18 10:07:45.124 AM [APP/1-7/bad-task] OUT
4/25/18 10:07:45.124 AM [APP/1-7/bad-task] OUT /hana/shared/ABC/xs/app_work...
4/25/18 10:07:45.127 AM [APP/1-7/bad-task] ERR 'non-existing...
4/25/18 10:07:45.127 AM [APP/1-7/bad-task] ERR operable program or batch fi...
4/25/18 10:07:49.043 AM [API] ERR Crashed instance [state CRASHED, index 7]...
4/25/18 10:08:22.774 AM [API] OUT Starting new task instance '09d52632-d616...
4/25/18 10:08:24.014 AM [APP/1-8/my-task] OUT
4/25/18 10:08:24.014 AM [APP/1-8/my-task] OUT /hana/shared/ABC/app_working/...
4/25/18 10:08:24.018 AM [APP/1-8/my-task] OUT Hello World
```

The log output reveals details of the failed task, which in the example above concerns the use of an unrecognized command (`non-existing`). The log lines relevant for a given task can be found by looking for the task name and task index within the log line **source** string, which is enclosed in square brackets in the log output, for example, `[APP/1-7/bad-task]`. The name and index associated with a task are displayed in the output of the `xs tasks` command.

The successful "Hello World" task shown in the log output has the index "8" and the name "my-task" and it was executed using the first droplet of the `my-app` application. As a result, the output of this task is marked with the source string `[APP/1-8/my-task]`.

Related Information

[Maintaining Applications in XS Advanced \[page 1684\]](#)

[The XS Command-Line Interface \[page 1653\]](#)

12.2.2.1.22 Scanning and Reporting XS Advanced Application Artifacts

Generate and display reports with details of dependencies between application artifacts and resources in XS advanced.

Web applications typically include a large number of software components such as Java archives (`.jar`) or Node.js modules (`.js`). In most cases, these artifacts are third-party (probably open-source) components that need special attention. And these components are not only in the application bundle pushed to the server; a significant amount of additional artifacts can be added by the platform itself during application deployment, for example: UI-framework artifacts, middleware, or low-level modules such as database drivers or cryptographic libraries. Regardless of origin, all artifacts in a staged application potentially contribute to productive code. For this reason, it is crucial to be able to get a complete overview of all artifacts that an application references either directly or indirectly, in order to identify the following:

- Vulnerable (open-source) components
- Components with well-known functional flaws
- Outdated or redundant component versions
- Components with license restrictions

The xs CLI commands `find-app-artifacts`, `find-droplet-artifacts`, `find-runtime-artifacts`, and `find-buildpack-artifacts` are suitable for analyzing artifact dependencies of a single XS advanced controller resource such as an application, a droplet (a staged application), a deployed run time, or a build pack. When scanning a broader range of resources on the server, for example, a whole space or even an organization, it makes sense to start with the command `xs find-artifacts`, which provides the same basic features for analyzing and filtering artifacts, but operates on a larger set of resources. In their initial version, all artifact-reporting tools for XS advanced applications are limited to the detection of Java archives and Node.js modules.

→ Tip

You can use operating-specific tools to pipe the report output generated by the `find-*-artifacts` commands to other commands such as `grep` or to a specific file. If the search completes successfully but nothing is found that matches the search query, the commands have process exit code "0". If the search query finds a match, "-1" is returned, and "1" if the commands finish unsuccessfully.

Displaying the Usage of Application Modules

After deploying an application `my-web-app` to the XS advanced platform, you can find out which version of a particular Node.js module (for example, `express`) was uploaded to the server during application deployment.

To scan the uploaded application's artifacts, run the command `xs find-app-artifacts` as a user with the privileges required to read application artifacts:

Note

The current `xs` target must be the space where the application is running.

Output Code

```
$ xs find-app-artifacts my-web-app -n express*

Finding artifacts of app "my-web-app" with name wild card "express*"...

+ APP "my-web-app" in space "SAP" of org "test" created at Jan 18, 2018
  12:33:03 PM (STARTED)
    NPM express 4.15.3
      path: /node_modules/express
    NPM express-session 1.15.3
      path: /node_modules/express-session

Found Artifacts   Affected Apps
-----
2                 1
```

The report returned by the `xs find-app-artifacts` command shows that the deployed application `my-web-app` uses the modules `express` version 4.15.3 and `express-session` version 1.15.3, and that both modules are located in the `node-modules` folder.

Tip

The command `xs find-app-artifacts <APPNAME>` displays **all** components uploaded with the application.

You can use the following options with **all** find-artifact commands:

- `--artifact-info`
Displays meta information about the artifacts, if available, for example, vendor, license, etc.
- `--artifact-tree`
Displays the logical dependency tree of the components, if available. This helps understand which artifacts are referenced transitively by a direct dependency.

Showing Which Artifacts the Platform Adds During Application Staging

`my-spring-app`) to the XS advanced platform, it is important to understand that, in addition to the Java archives (If you push a spring-based Java application (for example, `.jar` files) required to implement the application logic, the Java build pack also adds several Java archives from the build pack content itself, as well as from several run times. You can use the command `xs find-droplet-artifacts` to display a report with details of all the artifacts in the target application's current droplet. The report displayed by the `xs find-`

droplet-artifacts command also shows higher-level resources such as the application itself and the build pack and run times used, as shown in the following example:

Output Code

```
$ xs find-droplet-artifacts my-spring-app

Finding artifacts of app "my-spring-app"...

+ DROPLET of application "my-spring-app" with index 3 in space "PROD" of
org "test" created at Jan 19, 2018 1:47:39 PM (1 RUNNING, 1 STOPPED)
+ APP "my-spring-app" in space "PROD" of org "test" created at Sep 17, 2018
3:11:07 PM (STARTED)
  JAR spring-security-core 4.2.3.RELEASE
    path: /app/WEB-INF/lib/spring-security-core-4.2.3.RELEASE.jar
  JAR spring-security-oauth2 2.3.3.RELEASE
    path: /app/WEB-INF/lib/spring-security-oauth2-2.3.3.RELEASE.jar
  JAR spring-security-web 4.2.3.RELEASE
    path: /app/WEB-INF/lib/spring-security-web-4.2.3.RELEASE.jar
  JAR spring-web 4.3.15.RELEASE
    path: /app/WEB-INF/lib/spring-web-4.3.15.RELEASE.jar
  [...]
+ BUILDPACK "sap_java_buildpack" 1.7.8
  JAR httpclient 4.5.3
    path: /app/META-INF/.sap_java_buildpack/tomcat/impl/httpclient-4.5.3...
  JAR httpcore 4.4.6
    path: /app/META-INF/.sap_java_buildpack/tomcat/impl/httpcore-4.4.6.jar
  JAR java-container-security 0.30.1
    path: /app/META-INF/.sap_java_buildpack/tomcat/impl/java-container-sec
  [...]
+ RUNTIME "hanajdbc2" 2.3.53
  JAR ngdbc 2.3.53
    path: /app/META-INF/.sap_java_buildpack/hana_jdbc/ngdbc-2.3.53.jar
+ RUNTIME "sapjvm8_jre" 8.1.43
  JAR rt 1.8.0_181
    path: /app/META-INF/.sap_java_buildpack/sapjvm/lib/rt.jar
  [...]
+ RUNTIME "tomcat8" 8.5.32
  JAR tomcat-api 8.5.32
    path: /app/META-INF/.sap_java_buildpack/tomcat/lib/tomcat-api.jar
  JAR tomcat-coyote 8.5.32
    path: /app/META-INF/.sap_java_buildpack/tomcat/lib/tomcat-coyote.jar
  JAR tomcat-websocket 8.5.32
    path: /app/META-INF/.sap_java_buildpack/tomcat/lib/tomcat-websocket.jar
  [...]

Found Artifacts    Affected Droplets
-----
226                1
```

The example above shows that the spring-based application droplet contains the following artifacts:

- Application "my-spring-app"
- Build pack "sap-java-buildpack" (version 1.7.8)
- Run time "hanajdbc2" (version 2.3.53)
- Run time "sapjvm8_jre" (version 8.1.43)
- Run time "tomcat8" (version 8.5.32)

In total, the report found 226 artifacts matching the search criteria my-spring-app.

Listing the Artifact Versions Used in Productive Code

You can use the `xs` command-line tools to check if specific versions of an application component are still in use. This type of report is useful if a particular version of an application component is known to have a vulnerability or a proven functional flaw. For example, the `xs find-artifacts` command enables you to scan a large set of application resources on the server within a customizable search range: applications, droplets, run times, or build packs.

Since, generally speaking, the scope of the scan is determined by the privileges granted to the XS advanced platform user who runs the scan, it is recommended that application-artifact scans are run by a platform administrator user, in order to ensure that all resources available on the server are analyzed. You can limit the scope of a scan, for example, to search within a specific organization or space, by using the options `"-o <ORG>"` and `"-s <SPACE>"`.

The following example shows how to run a search with the scope set to the Java archive `security-commons` with major version "0" and minor version "26", which is outdated and needs to be replaced:

Output Code

```
$ xs find-artifacts -n security-commons* -v 0.26.*
Finding all artifacts of droplets with name wild card "security-commons*"
and version wild card "0.26.*"...
+ DROPLET of application "my-java-app" with index 1 in space "PROD"
  of org "test" created at Jan 20, 2018 9:02:42 AM (1 RUNNING, 1 STOPPED)
  + BUILDPACK "sap_java_buildpack" <outdated version, latest version:1.7.8>...
    JAR security-commons 0.26.13 (OUTDATED)
    path: /app/META-INF/.sap_java_buildpack/tomcat/impl/security-commons-0.26
Found Artifacts      Outdated Artifacts      Affected Droplets
-----
2                    2                    1
There are outdated artifact versions in search scope.
Restage and restart affected apps to replace them with current versions
available on the platform.
```

The command output in the example above shows that a droplet of a running instance of the application "my-java-app" is using an outdated version (0.26.13); the dependency was apparently introduced by "sap_java_buildpack". The first "OUTDATED" tag in the example above indicates that the build pack "sap_java_buildpack" on the platform has been updated to version 1.7.8 after the application "my-java-app" was staged; restaging the application should update this dependency.

Note

After staging the `my-java-app` application with the current build pack, it is necessary to restart the application to activate the updates.

Related Information

[The XS Command-Line Interface \[page 1653\]](#)

12.2.2.2 Maintaining XS Advanced Run-Time Instances with the XSA CLI

Manage XS advanced services without stopping the SAP HANA database.

Similar to the `HDB` command that lists all SAP HANA services, the `xsa` command enables you to manage XS advanced instances without having to stop and restart the SAP HANA database. With the `xsa` command, you can manage XS advanced components, for example, the controller (`xscontroller`), the execution agent (`xsexecutionagent`), and the `xsaaserver` that provides services for User Accounts and Authentication (UAA).

With the `xsa` command, you can perform the following actions on an instance of the XS advanced run time:

- Enable
- Disable
- Restart
- Set a domain certificate
- Reset the default domain certificate for an XS advanced run-time instance
- Restage all applications that are running when the XS advanced run-time instance is restarted
- Backup and restore the Secure Store file system
- Delete user data from log files

Related Information

[The XSA Command Reference \[page 1740\]](#)

12.2.2.2.1 The XSA Command Reference

Commands to help maintain XS advanced run-time instances, for example, enable, disable, or restart.

```
XSA {COMMAND} [--OPTIONS]
```

The following table lists the `xsa` commands that are most commonly used:

! Restriction

To use the `xsa` command, you must log on as the operating-system user `<SID>adm`.

XSA Commands Overview

Command	Description
<code>help</code>	Display an overview of all command help parameters and options
<code>enable</code>	Enable a disabled XS advanced run-time instance

Command	Description
<code>disable</code>	Disable a running instance of the XS advanced run time
<code>restart</code>	Disable and re-enable an XS advanced run-time instance
<code>set-certificate</code>	Set the certificate for the default domain after the XS Controller is shut down
<code>reset-certificate</code>	Reset the certificate for the default domain and restart the XS advanced run-time instance
<code>trust-certificate</code>	Add a trusted certificate
<code>restage-at-startup</code>	Restage all applications that are running when the XS advanced instance (controller) is restarted
<code>backup-ssfs</code>	Store the current XS advanced Secure Store File System (SSFS) configuration in an encrypted database table
<code>recover-ssfs</code>	Restore the XS advanced Secure Store File System (SSFS) configuration from an encrypted database table
<code>diagnose</code>	Run tests that help diagnose problems with the XS advanced system
<code>backup-fss</code>	Create a backup copy of the file-system services in the database
<code>recover-fss</code>	Restore the file-system service instances from the database
<code>collect-traces</code>	Creates a zip containing the trace files of the XS advanced services
<code>delete-personal-data</code>	Erase user data from the log files
<code>du</code>	Provides the disk usage for application instances and file-system service
<code>grant-privileges-to-support-user</code>	Grants an SAP HANA user access to XS advanced database schemas
<code>list-tenants</code>	Lists all (tenant) databases registered for XS advanced
	<div style="background-color: #f0f0f0; padding: 10px; border-left: 2px solid #0070c0;"> <p>→ Tip</p> <p>You can also verify the credentials of the technical XS advanced users within tenants used by XS advanced.</p> </div>
<code>select-xsa-runtime-db</code>	Activates XS advanced for the specified tenant database

enable

Enable a disabled XS advanced run-time instance.

Usage

```
XSA enable [--localhost] [--verbose] [--async]
```

Options

Command Options Overview

Option	Description
--localhost	Restrict the operation to the local host
--verbose	Display all available information during the selected operation
--async	Return to the command prompt without waiting for the command to complete

disable

Disable a running instance of the XS advanced run time.

Usage

```
XSA disable [--localhost] [--verbose] [--async]
```

Options

Command Options Overview

Option	Description
--localhost	Restrict the operation to the local host
--verbose	Display all available information during the selected operation
--async	Return to the command prompt without waiting for the command to complete

restart

Disable and re-enable an XS advanced run-time instance.

Usage

```
XSA restart [--localhost] [--verbose] [--async]
```

Options

Command Options Overview

Option	Description
<code>--localhost</code>	Restrict the operation to the local host
<code>--verbose</code>	Display all available information during the selected operation
<code>--async</code>	Return to the command prompt without waiting for the command to complete

set-certificate

Set the certificate for the default domain of an XS advanced run-time instance.

i Note

The `set-certificate` command enables you to set a valid domain certificate without the need to start the XS advanced controller first.

By default, all applications and the XS advanced controller are reached by means of the default domain that is configured during the installation operation. To enable HTTPS communication, it is necessary to configure an SSL certificate for the default domain or, alternatively, use the self-signed certificate provided by XS advanced. If the certificate provided for the default domain expires, neither the applications nor the XS advanced controller can be reached. In addition, the XS Controller will not be reachable after a restart because the initial availability checks will fail. The `set-certificate` command enables you to resolve the problem by setting the certificate for the default domain.

Usage

```
XSA set-certificate [--cert <CERTIFICATE_FILE>] [--key <KEY_FILE>] [--verbose]
[--async]
```

Options

Command Options Overview

Option	Description
<code>--cert</code> <code><CERTIFICATE_FILE></code>	Specify the file containing the public certificate to use
<code>--key <KEY_FILE></code>	Specify the file containing the private key to use
<code>--verbose</code>	Display all available information during the selected operation
<code>--async</code>	Return to the command prompt without waiting for the command to complete

reset-certificate

Reset the certificate for the default domain and restart the XS advanced run-time instance.

→ Tip

The `reset-certificate` command can only be used to generate a new **self-signed certificate**, which is helpful in cases when you just want to get the XS advanced controller running again after a certificate is outdated.

By default, all applications and the XS advanced controller are reached by means of the default domain that is configured during the installation operation. To enable HTTPS communication, it is necessary to configure an SSL certificate for the default domain or, alternatively, use the self-signed certificate provided by XS advanced. If the certificate provided for the default domain expires, neither the applications nor the XS advanced controller can be reached. In addition, the XS Controller will not be reachable after a restart because the initial availability checks will fail. The `reset-certificate` command enables you to reset the certificate for the default domain to resolve the problem.

Usage

```
XSA reset-certificate [--verbose] [--async]
```

Options

Command Options Overview

Option	Description
<code>--verbose</code>	Display all available information during the selected operation
<code>--async</code>	Return to the command prompt without waiting for the command to complete

trust-certificate

Adds a trusted certificate, even if the XS advanced controller is not running.

Usage

```
XSA trust-certificate [--alias <ALIAS>] [--cert <CERTIFICATE_FILE>] [--hana-broker] [--client-auth] [--verbose]
```

Options

Command Options Overview

Option	Description
<code>-a, --alias <ALIAS></code>	The alias of the trusted certificate

Option	Description
<code>-c, --cert <CERTIFICATE_FILE></code>	The path to the X.509 certificate file in PEM format.
<code>--client-auth</code>	Use the certificate for validation of certificate-based client authentication
<code>--hana-broker</code>	Add the certificate to the service bindings issued by the SAP HANA service broker.
<code>--verbose</code>	Display all available information during the selected operation

restage-at-startup

On restart of the XS advanced instance (controller), restage all applications that were running in the XS advanced instance before the restart. Restaging applications ensures that they have access to any security fixes for deployment build packs and run-time environments that are available.

Usage

```
XSA restage-at-startup [--verbose]
```

→ Tip

The XS advanced instance is not restarted automatically; you must restart the XS advanced instance yourself manually.

Options

Command Options Overview

Option	Description
<code>--verbose</code>	Display all available information during the selected operation

backup-ssfs

Store the current XS advanced Secure Store File System (SSFS) configuration in an encrypted database table.

Usage

```
XSA backup-ssfs [--verbose]
```

i Note

The `backup-ssfs` command replaces the command `save-ssfs-to-dbss`.

Options

Command Options Overview

Option	Description
<code>--verbose</code>	Display all available information during the selected operation

recover-ssfs

Restore the XS advanced Secure Store File System (SSFS) configuration from an encrypted database table.

Usage

```
XSA recover-ssfs [-u <system_db_user>] [-p <system_db_password>] [-n <db_host_port>] [--verbose]
```

i Note

The `recover-ssfs` command replaces the command `restore-dbss-to-ssfs`.

Options

Command Options Overview

Option	Description
<code>-u</code> <code><system_db_user></code>	The name of the system database user
<code>-p</code> <code><system_db_password></code>	The password of the system database user. i Note If you do not provide the password with the command, you are prompted to provide it interactively.
<code>-n <db_host_port></code>	The <code><server>[:<port>]</code> to use to establish a connection to the tenant database where XS advanced keeps its persistence, for example, in a multi-tenant database container (MDC) system that was migrated from a single database container system. i Note This information is not required if XS advanced keeps its persistence in the System database (default).
<code>--verbose</code>	Display all available information during the selected operation

diagnose

Run tests that help diagnose problems with the XS advanced system. The command writes its findings to `stdout` as well as to a dedicated trace file. The path to the location of the trace file is added to the end of the output written to `stdout`.

Usage

```
XSA diagnose [--verbose]
```

backup-fss

Creates a backup of the file-system services in the database.

Usage

```
XSA backup-fss [--verbose]
```

recover-fss

Restores the file-system service instances from the database.

Usage

```
XSA recover-fss [--verbose]
```

delete-personal-data

Remove user data from the log files generated in an XS advanced run-time instance. The `delete-personal-data` command triggers log-rotation for all log files and removes all user-related data (for example, user names, IP addresses) from the log files up to a specified point in time. You can use the `delete-personal-data` command to clean up the following types of XS advanced log files:

- XS advanced Controller log files (`controller_#.log`)
- XS advanced Execution Agent log files (`ea_#.log`)
- HTTP application-access files (`webdispatcher`)
- Application trace files (`myApp.trace`)
- User Account and Authentication (UAA) service artifacts

Usage

```
XSA delete-personal-data [--until <DATE>] [--force] [--exclude <EXCLUDES>] [--user <NAME>] [--verbose]
```

Options

Command Options Overview

Option	Description
<code>-u, --until <DATE></code>	The date that specifies the last day of the period of time for which you want to delete log-file entries that contain user-related information.
	→ Tip <code><DATE></code> must be specified in the format YYYY-MM-DD, for example, 2017-06-30.
<code>-f, --force</code>	Force execution of the data-deletion command without requiring any user confirmation
<code>-e, --exclude <EXCLUDES></code>	Exclude the specified log types from the deletion operation. Multiple log types are specified in a comma-separated list consisting of the following values: [PLATFORM, AUDIT, APP, ACCESS, USERS]
<code>--user <NAME></code>	The name of the XS advanced administrator user in whose account the data-deletion command will run
<code>--verbose</code>	Display all available information during the selected operation

collect-traces

Creates a Zip file containing the trace files written by the XS advanced services; the created zip file overwrites any existing Zip file previously created by the `collect-traces` command.

Usage

```
XSA collect-traces --output <xsa_traces.zip>
```

Options

Command Options Overview

Option	Description
<code>-o, --output <xsa_traces.zip></code>	The path to the Zip file containing the collected trace files

grant-privileges-to-support-user

Grants an SAP HANA user access to XS advanced database schemas. If the target user does not already exist, a new user is created with the specified name.

Usage

```
XSA grant-privileges-to-support-user --support-user <SUPPORT_USER_NAME>
[--system-user <SYSTEM_USER_NAME>] [--system-password <SYSTEM_USER_PASSWORD>]
[--support-password <SUPPORT_USER_PASSWORD>] [--grant-all]
```

Options

Command Options Overview

Option	Description
<code>--system-user</code> <code><SYSTEM_USER_NAME></code>	The name of the <code>SYSTEM</code> user of the XS advanced run-time database
<code>--system-password</code> <code><SYSTEM_USER_PASSWORD></code>	The password of the <code>SYSTEM</code> user of the XS advanced run time database. ⚠ Caution For security reasons, it is not recommended to use the <code>--password</code> option to specify a password for the <code>SYSTEM</code> user. If no password is specified, the system interactively requests one.
<code>-u, --support-user</code> <code><SUPPORT_USER_NAME></code>	The name of the support user to whom support privileges should be granted. i Note If the target user does not already exist, a new user is created with the specified name.
<code>-p, --support-password</code> <code><SUPPORT_USER_PASSWORD></code>	The password of the support user to whom access privileges are to be granted. i Note If no password is specified, the system interactively requests one, and creates a new user with the specified name, if the target user does not already exist.
<code>-a, --grant-all</code>	Grant all access privileges, not only <code>SELECT</code> privileges

du

Provides the disk usage for application instances and file-system service.

Usage

```
XSA du [--app] [--fss] [--orderby <ORDERBY>] [--orderdir <ORDERDIR>]
```

Options

Command Options Overview

Option	Description
<code>-a, --app</code>	Display the disk usage for application instances running on target XS advanced hosts
<code>-f, --fss</code>	Display the disk usage for the file-system service
	i Note Requires SAP HANA to be started.
<code>--orderby <ORDERBY></code>	Sort by specific columns for application disk usage
	→ Tip The option accepts comma-separated values.
<code>--orderdir <ORDERDIR></code>	Specify the sort order direction for application disk usage, for example, asc(ending) or desc(ending)

list-tenants

Display details of all (tenant) databases registered for use with XS advanced model. It is also possible to check the validity of the user credentials that XS advanced uses to connect to each registered database.

Usage

```
list-tenants [--user <XSA admin user>] [--password <XSA admin password>] [--check-credentials]
```

Options

Command Options Overview

Option	Description
<code>-u, --user <XSA admin user></code>	The name of an XS advanced administration user (for example, XSA_ADMIN)
<code>-p, --password <XSA admin password></code>	The password for the XS advanced administration user specified with the --user option
<code>--check-credentials</code>	Verify the credentials of the XS advanced technical users that XS advanced uses to connect to each registered database

select-xsa-runtime-db

Activates XS advanced for the specified tenant database.

Usage

```
select-xsa-runtime-db --tenant-db-name <tenant_db_name>
[--tenant-db-system-user <tenant_db_system_user>] [--tenant-db-system-user-
password <tenant_db_system_user_password>]
[--system-db-password <system_db_password>] [--system-db-user <system_db_user>]
```

Options

Command Options Overview

Option	Description
-n, --tenant-db-name <tenant_db_name>	The name of the tenant database
-t, --tenant-db-system-user <tenant_db_system_user>	The name of the tenant database's SYSTEM user
-p, --tenant-db-system-user-password <tenant_db_system_user_password>	The password of the tenant database's SYSTEM user ⚠ Caution For security reasons, it is not recommended to use the --password option to specify a password for the SYSTEM user. If no password is specified, the system interactively requests one.
-c, --system-db-password <system_db_password>	The password of the SYSTEM database user ⚠ Caution For security reasons, it is not recommended to use the --password option to specify a password for the SYSTEM user. If no password is specified, the system interactively requests one.
-s, --system-db-user <system_db_user>	The name of the SYSTEM database user

Related Information

[Maintaining XS Advanced Run-Time Instances with the XSA CLI \[page 1740\]](#)

12.2.2.2.2 Data Protection and Privacy Tools in XS Advanced

Ensure the protection of private user data in XS advanced.

To comply with data privacy regulations, it must be possible to delete personal and private data stored in an XS advanced system. The XS advanced run-time environment stores the following types of user-related data:

- IP addresses and user names in the audit log and trace files
- IP addresses in the router access logs
- User names in the application log files
- Shadow users in the UAA

XS advanced administrators can use the command `XSA delete-personal-data` to remove personal and private data. For example, the following command ensures the deletion of all personal data until the second of January 2018:

```
XSA delete-personal-data -u 2018-01-02
```

The `--exclude` enables you to prevent certain types of user information from being included in the delete operation, for example, if it is necessary to keep the information in order to meet other regulations, such as maintaining the integrity of the audit log, where the information is required. The `delete-personal-data` command enables you to delete sensitive information from the following types of sources:

- PLATFORM
The platform logs including the event log in the database and the trace files
- AUDIT
The audit log files of XS advanced
- APP
The application log files
- ACCESS
The Webdispatcher access logs
- USERS
The shadow users in the User Account and Authentication (UAA) server

The following command deletes personal data from all sources in XS advanced **except** the Audit log up until February 2, 2018:

```
XSA delete-personal-data -u 2018-01-02 --exclude AUDIT
```

⚠ Caution

After deletion data cannot be recovered.

Data cannot be recovered after deletion. The only exception to this rule is the shadow users in the User Account and Authentication server (UAA), which are created again the next time the users log in. Log files are not included in the SAP HANA backup.

Related Information

[The XSA Command Reference \[page 1740\]](#)

12.2.3 Maintaining the XS Advanced Run-time Environment with a Graphical User Interface

Use a graphical user interface to administrate and maintain XS advanced-model run-time components.

SAP HANA XS advanced model includes a Web-browser-based administration tool called XS Advanced Cockpit with a graphical user interface that enables you to maintain important elements of the XS advanced application-development environment, for example, security and authentication methods.

The XS Advanced Cockpit includes a selection of tools and features which enable you to configure and maintain the basic administration-related elements of the application-development process for the XS advanced run-time environment.

Related Information

[Maintaining the XS Advanced Runtime Environment with SAP HANA XS Advanced Cockpit \[page 1753\]](#)

[Maintaining the XS Advanced Run-time Environment with a Command-Line Interface \[page 1652\]](#)

[Scheduling Jobs in XS Advanced \[page 1790\]](#)

12.2.3.1 Maintaining the XS Advanced Runtime Environment with SAP HANA XS Advanced Cockpit

SAP HANA XS Advanced Cockpit is a graphical user interface that enables you to configure and maintain the basic administration-related elements of the application development process for the XS advanced runtime environment. The home page appears after you log on to XS Advanced Cockpit. The navigation pane on the left contains various runtime options. You can find more runtime options under *More...* in the left navigation pane. Based on your role in an organization, you can see one or more options for performing tasks specific to your role. The pane on the right displays organizations. You can be a member of one or more organizations. To work with applications and services, you need to navigate to spaces within an organization. For more information about working with organizations and spaces, see *Maintaining Organizations and Spaces in XS Advanced* in the *Related Information* section. The table below provides a brief overview and scope of the different runtime options:

Runtime Options	Description	Scope
Organization and Space Management	Create, list, or delete organizations and spaces in the XS advanced model runtime.	<ul style="list-style-type: none"> Managing organizations Managing spaces Managing users in organizations and spaces Managing XS advanced business-user roles for organizations and spaces Managing applications and services within a space
SAML Identity Provider	Configure an SAML identity provider for use by XS advanced applications that need to authenticate the XS advanced business users signing in by means of SSO.	<ul style="list-style-type: none"> Managing SAML identity providers, including IDP metadata, certificates, and destinations
Tenant Databases	Manage, maintain, and configure SAP HANA tenant databases for use with SAP HANA XS advanced model applications.	<ul style="list-style-type: none"> Managing tenant databases in XS advanced Creating tenant databases in XS advanced Preparing tenant databases
Host Management	Manage the SAP HANA hosts that are pinned to SAP HANA XS advanced applications or spaces.	<ul style="list-style-type: none"> Displaying a list of SAP HANA hosts that are pinned to XS advanced applications or spaces
User Management	Maintain and manage database users for SAP HANA XS advanced	<ul style="list-style-type: none"> Creating XS advanced platform users
Application Monitoring	Monitor the system resources used by applications running in the XS advanced model runtime environment.	<ul style="list-style-type: none"> Monitoring application use of system resources
Trusted Certificates	Manage and maintain the certificates used to establish secure and trusted connections between SAP HANA systems and SAP HANA XS advanced applications.	<ul style="list-style-type: none"> Adding certificates Displaying a list of certificates and details
Application Management	Perform various actions on an application, such as scaling an application, stopping or restarting an application, establishing secured access, verifying application logs or events for troubleshooting.	<ul style="list-style-type: none"> Scaling application instances up or down Managing application security Troubleshooting an application

Related Information

[Maintaining Organizations and Spaces in XS Advanced \[page 1758\]](#)

12.2.3.1.1 Managing Users in XS Advanced

The *User Management* option in the XS Advanced Cockpit enables you to create a new user or add a user from SAP HANA so that users have the required authentication and authorization to work with the various runtime options included with the XS Advanced platform. You can perform the following tasks:

- Create a new user
- Modify user details
- Search for a user
- Delete a user

12.2.3.1.1.1 Create Users in XS Advanced

Context

This option enables you to create users or to promote existing SAP HANA users to access organizations and spaces in the SAP HANA XS Advanced runtime (including XS Advanced Cockpit).

Procedure

1. In the home navigation pane, choose *User Management*.
2. On the *User Management* screen, choose any of the following options:
 - *New User*: Choose this option to create a user.
 - *Migrate SAP HANA User*: Choose this option to promote an existing SAP HANA user to an XS Advanced runtime user. You can either provide the SAP HANA user ID or select the user from the user list.

Related Information

[User Details \[page 1757\]](#)

12.2.3.1.1.2 Manage Users

Context

You can perform multiple actions to maintain user details, such as update user details, reset password, assign role collection, search, or delete users from the XS Advanced Cockpit.

Procedure

1. In the home navigation pane, choose *User Management*.
2. On the *User Management* screen, perform any of the following tasks:

Task	Choose	Description
Update user details	 (Edit User)	You can modify user details such as first name, last name, or e-mail address for the selected user. This option is available underneath the <i>Actions</i> column.
Reset password of a user	 (Change Password)	You can change the password required to access the XS Advanced runtime including the XS Advanced Cockpit. This option is available underneath the <i>Actions</i> column.
Assign role collections to a user	 (Assign Role Collections)	As well as the default role collection that is available for all users, you can add additional role collections and assign them to specific users. This option is available underneath the <i>Actions</i> column.
Delete a user	 (Delete User)	This option deletes a user. This option is available underneath the <i>Actions</i> column.

Task	Choose	Description
Search for a specific user	Search	If your member management screen contains a long list of users, you can use the search option to list a specific user. You can search for users using any of the following user details: member ID, user name, or role collection.

Related Information

[Assign Roles to Role Collection \[page 1781\]](#)

[Control Access to Applications \[page 1769\]](#)

[User Details \[page 1757\]](#)

12.2.3.1.1.2.1 User Details

You can view the details of an XS Advanced platform user with the XS Advanced Cockpit. These are the details that are provided in the *New User* dialog.

User Interface Element	Description
User ID	Unique ID of the XS Advanced user. This is a mandatory field.
First Name	First name of the user.
Last Name	Last name of the user.
E-Mail	E-mail address of the user to whom the account belongs. The e-mail address must be unique to the user. This is a mandatory field.
Password	Password for the XS Advanced user. This is a mandatory field.

12.2.3.1.2 Maintaining Organizations and Spaces in XS Advanced

Organization is a logical entity to group spaces. By default, an organization does contain a space. You can create one or more spaces to enable developers to collaborate by sharing resource, services, and applications. Access to the shared resources, services, and applications is controlled by roles, for example, Organization Manager or Organization Auditor. The role defines the scope of permissions available for a user in an organization. For example, an Organization Manager can add new users to organizations, create, modify, or delete organizational spaces, and add domains to the organization.

A space enables users to develop and maintain applications. Each space provides shared access to users of the space for application development, deployment, and maintenance. Access to the resources is controlled by roles, for example, Space Manager, Space Developer, or Space Auditor. The role defines the scope of permissions available for a user in a space. For example, a Space Developer can deploy and start an application.

You can perform the following tasks within an organization and space:

- Create an organization
- Manage an organization
- Maintain users in an organization
- Create a space
- Manage a space
- Maintain users in a space

Related Information

[Maintaining Organizations and Spaces in XS Advanced \[page 1663\]](#)

12.2.3.1.2.1 Create an Organization

Context

You create an organization to establish a collaborative environment for sharing resources, services, and applications.

Procedure

1. In the home navigation pane, choose [Organizations](#).
2. On the [Organizations](#) screen, choose [New Organization](#).
3. Provide the following details:

User Interface Element	Description
Name	Provide a name for the organization. This field is mandatory.
Assign Roles to User	This option provides the self-service option of assigning one or both of the roles (Organization Manager or Organization Auditor). You need to have these roles to perform organizational tasks.

12.2.3.1.2.2 Manage an Organization

Context

You can update the details or delete organizations.

Procedure

1. In the home navigation pane, choose [Organizations](#).
2. On the [Organizations](#) screen, perform any of the following tasks:

Tasks	Choose	Description
Update organizational details	 (Edit)	You can update details such as the name of an organization.
Delete an organization	 (Delete)	You can remove an organization.

Tasks	Choose	Description
Search for an organization	Search	If you want to view only a specific organization from the list of organizations, enter the organization name in the search field.

12.2.3.1.2.3 Maintain Users in an Organization

Context

To perform tasks within an organization, you need to add users to specific roles within the organization. Roles give controlled access to users within an organization.

Procedure

1. In the home navigation pane, choose *Organizations*.
2. In the right pane, choose an organization tile.
3. In the members navigation pane, choose *Members*.
4. On the *Members* page, choose *Add Members*.
5. Enter one or more user IDs.
6. Assign one or both of the roles.

Related Information

[Maintaining Platform Users in XS Advanced \[page 1671\]](#)

12.2.3.1.2.3.1 Organizational Roles

You can grant or restrict access to organizations by assigning roles. The following table lists the roles that you can assign to XS Advanced users in an organization:

Role	Description
Organization Manager	An Organization Manager can perform the following tasks: <ul style="list-style-type: none"> • Create and manage organization users • Create, modify, or delete organizational spaces • Add domains to the organization
Organization Auditor	An Organization Auditor can perform the following tasks: <ul style="list-style-type: none"> • View all users in the organization • View the roles assigned to a user (or users) in the organization • View spaces within an organization

12.2.3.1.2.4 Create a Space

Prerequisites

You are already a member of an organization and have the Organization Manager role.

Context

You create a space to enable users to develop and maintain applications.

Procedure

1. In the home navigation pane, choose [Organizations](#).
2. On the [Organizations](#) screen, select an organization.
3. On the [Spaces](#) screen, choose [New Space](#).
4. Provide the following details:

User Interface Element	Description
Space Name	Provide a name for the space. This field is mandatory.

User Interface Element	Description
Assign space roles	This is a self-service option for assigning one or all of the roles (Manager, Developer, or Auditor). You need to have these roles to perform space-related tasks.

12.2.3.1.2.5 Manage a Space

Prerequisites

You have already selected an organization.

Context

A space contains various deployed applications and resources such as services and members. You can perform various tasks in a space, such as stop, start, delete, or search for an application. You can also do the following:

- Bind a service to an application deployed to the space
- Assign (pin) one or more hosts (SAP HANA systems) to a space

Procedure

1. In the space navigation pane, choose [Spaces](#).
2. On the [Spaces](#) screen, perform any of the following tasks:

Task	Choose	Description
Start a stopped application	 (Start)	You can use this option to start an application from its stopped state.
Stop an application that is running	 (Stop)	You can use this option to stop an application that is currently running.

Task	Choose	Description
Delete an application	 (Delete)	You can use this option to remove an application from a space.
Search for an application	Search	If a space contains a long list of deployed applications, you can use the search option to list a specific application. Search for an application using the application name.
Manage an application	application name	You can use this option to perform multiple tasks within an application, such as define security, check application logs or events, and so on.
Assign one or more hosts to a space	Pinned Hosts	<p>You can use this option to select one or more hosts from the displayed list. The applications in the space will be deployed to the selected hosts. You can define one of the following modes:</p> <ul style="list-style-type: none"> ○ Strict: In this mode, an application cannot start if the pinned host is not available. ○ Relaxed: In this mode, an application can start on another host if the pinned host is not available.
<div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> <p>i Note</p> <p>You can switch from one mode to the other.</p> </div>		
Add users to a space	Members	You can use this option to grant permissions to users to perform specific tasks within the space. For more information, see Maintain Users in a Space [page 1764] .
Use a service within the space	Service Marketplace / Service Instances / User-Provided Services	You can use this option to use the services available within the space with any applications deployed to the space.

Related Information

[Managing XS Advanced Applications \[page 1766\]](#)

12.2.3.1.2.6 Maintain Users in a Space

Prerequisites

- You are already a member of the organization that contains the space.
- You have the Space Manager role in the space.

Context

To enable users to perform tasks within a space, you need to add the users to the space within the organization. Roles give controlled access to users within a space.

Procedure

1. In the space navigation pane, choose *Spaces*.
2. On the *Spaces* screen, select a space.
3. In the navigation pane, choose *Members*.
4. On the members page, choose *Add Members*.
5. Enter one or more user IDs.
6. Assign one or all of the roles.

Related Information

[Maintaining Platform Users in XS Advanced \[page 1671\]](#)

12.2.3.1.2.6.1 Space Roles

You can grant or restrict access to spaces by assigning roles. The following table lists the roles that you can assign to XS Advanced users in a space:

Role	Description
Space Manager	<p>A Space Manager can perform the following tasks:</p> <ul style="list-style-type: none"> • Manage users in the selected space • View details of applications running in the space (for example: status, instances, service bindings, and resource usage)
Space Developer	<p>A Space Developer can perform the following tasks:</p> <ul style="list-style-type: none"> • Deploy, start, stop an application • Bind an application to (or unbind an application from) a service • View details of applications running in the space (for example: status, instances, service bindings, and resource usage)
Space Auditor	<p>A Space Auditor can perform the following tasks:</p> <ul style="list-style-type: none"> • View details of applications running in the space (for example: status, instances, service bindings, and resource usage)

12.2.3.1.3 Monitoring the SAP HANA XS Advanced Model Runtime Environment

Monitoring provides details of system resources used by application instances running in the XS Advanced model runtime environment. You also have the option to sort and group resources. You can also monitor a specific application instance.

For example, you can see memory allocation and how long the application has been running.

12.2.3.1.3.1 View Utilized System Resources

Context

You can view the system resources utilized by an application.

Procedure

1. In the space navigation pane, choose [Spaces](#).
2. In the space navigation pane, choose [Monitoring](#).
Applications and their respective consumption details are listed.
3. To sort or group applications based on resource consumption details, choose  (View Settings).
4. In the [View](#) dialog, choose [Sort By](#) or [Group By](#) to arrange applications.

12.2.3.1.3.1.1 Resource Details

The Monitoring screen provides resource details of the XS applications running in the XS Advanced runtime. The following table contains description of the resource details:

User Interface Element	Description
Name	Name of the application
Memory (KB)	Amount of memory used by an application
CPU (ms)	Amount of processing time (in milliseconds) used by the application
User Mode (ms)	Amount of user-mode time (in milliseconds) consumed by an application. In user mode, an application cannot directly access hardware or modify memory; it can only do so by means of a proxy such as an API.
Kernel Mode (ms)	Amount of kernel-mode time (in milliseconds) consumed by an application. In kernel mode, an application has unrestricted access to CPU instructions and memory addresses.
Access Count	The number of times an application has been accessed.
MTA	Name of the multi-target application (MTA) to which the listed application belongs.
Host	The name of the host where the listed instance of an XS Advanced application is running.

12.2.3.1.4 Managing XS Advanced Applications

The application overview page provides various options available for the application.

You can stop an application that is currently running. You can also restart the application. For example, if you need to make changes to the active application, you can stop and restart the application to apply the changes.

You can also create a new application instance. For example, if the performance of your application is impaired because of excessive load on the application, you can manually start a new instance of the application to take up any new load.

12.2.3.1.4.1 Manage an Application

Context

You can use the application overview page to perform various tasks.

Procedure

1. In the spaces navigation pane, select an application.
2. On the [Overview](#) page, you can perform the following tasks:

Task	Choose	Description
Start an application that is currently stopped	Start	If an application is in a stopped state, use this option to start it again.
Stop an active application	Stop	If you want to stop an application, use this option.
Delete an application	Delete	This option deletes the application.
Create a new instance of an application	+Instance, - Instance	The + Instance option creates a new instance of an application. To delete an instance, choose - Instance.
View system resources consumed by an instance of the application	 Instance Monitoring	This option displays system resources consumed by an application instance. You can select an instance to view the resource consumption by each process within the instance.

Task	Choose	Description
View SAP HANA systems pinned to the application	 Pinned Hosts	<p>This option displays SAP HANA systems that are pinned to host your application. The <i>Pinned Hosts</i> screen displays the following options:</p> <ul style="list-style-type: none"> ○ Pin to Host: Choose this option to pin the application to the listed hosts. Strict is the default pin mode. In this mode, the application cannot start if the pinned host is not available. You can also choose the Relaxed pin mode. In this mode, the application can start on any other host if the pinned host is not available. ○ Change Pin Mode: Choose this option to switch between the available pin modes.
View logs specific to your application	 Logs	<p>This option displays logs generated by the application. You can choose the following options on the <i>Logs</i> screen:</p> <ul style="list-style-type: none"> ○ Change Log Level: Use this option to display log details of a specific type, such as fatal, error, warning, and so on. ○ Source: View logs of the specified source. ○ Type: View logs of specified types. ○ Recent /All: View recent or all logs. ○ Download: Download recent or all logs.
View actions performed on your application	 Events	This option displays actions performed on the application.
Bind an application to a service available in the service marketplace	Service Bindings	This option enables you to use a service available from the service marketplace with your application.

Related Information

[Managing Services in XS Advanced \[page 1787\]](#)

[Displaying Application Information in XS Advanced \[page 1676\]](#)

12.2.3.1.4.2 Control Access to Applications

Context

You can create user roles in the SAP HANA XS Advanced Cockpit to control access to XS Advanced applications. These roles are derived from role templates defined in the security description (xs-security.json) of applications that have been registered as OAuth 2.0 clients in the User Account and Authentication (UAA) service during application deployment. The application security description file also contains details of the authorization scopes that are used for application access, and defines any attributes that need to be applied. You can then add roles to role collections before the roles are assigned to SAP HANA database users or users logging on with SAML 2.0 assertions.

Procedure

1. On the organization overview screen, choose a space.
A list of applications deployed in the space appears.
2. In the application pane, select an application from the *Name* column.
The application overview page appears.
3. In the navigation pane, choose *Security*.
4. Perform the following actions to control secured access to the application:

Task	Choose	Description
Create a new role	Roles from the navigation pane	<p>This option enables you to create roles. A role is an instance of a role template. You can create a role based on a role template and assign the role to a role collection. Role collections are then assigned to SAP HANA users or SAML 2.0 groups. Provide the following details to create a role:</p> <ul style="list-style-type: none"> ○ Name: Name of the role ○ Description: Additional details about the role ○ Template: Templates are predefined for an application. Choosing a template identifies the scope and attribute applicable for the role. ○ Attribute: Provides attribute details applied to the role. Depending on the value of the attributes defined, access to resources is either granted or restricted. For example, in a sales scenario, the attribute region emea could be used to restrict access to the sales orders for the geographical region EMEA.
Update a role description	Roles from the navigation pane. In the <i>Roles</i> pane, choose  (Edit) underneath the <i>Actions</i> column.	This option enables you to update the role description.
Delete a role	Roles from the navigation pane. In the <i>Roles</i> pane, choose  (Delete) underneath the <i>Actions</i> column.	This option enables you to delete a role.
Assign roles to role collections	Roles from the navigation pane. In the <i>Roles</i> pane, choose  (Add to role collection) underneath the <i>Actions</i> column.	This option enables you to assign roles to role collections.
Check scopes applicable for the role	Scopes from the navigation pane	This option enables you to view permissions available for a role.

Task	Choose	Description
Check attributes	Attributes from the navigation pane	Attributes define information that comes with the respective user, for example, 'cost center' or 'country'. This information can only be resolved at runtime.
Check role templates available for the application	Role templates from the navigation pane	The role template defines the type of access permitted for an application, for example, the authorization scope and any attributes that need to be applied.

12.2.3.1.5 Managing Hosts in XS Advanced

This page displays the host systems on which XS Advanced applications are running. It provides specific details about the number of applications running on a host system and the spaces pinned to it. It also displays whether an application can only run on a specific host.

Related Information

[Maintaining Host Pinning \[page 1717\]](#)

12.2.3.1.5.1 View Host Systems

Context

Procedure

1. In the home navigation pane, choose *Host Management*.
2. Select a specific host to view the applications that can only run on this host and the spaces that are pinned to it.

12.2.3.1.5.1.1 Host System Details

The host management page displays the following information:

User Interface Element	Description
Host ID	Displays the host system.
Exclusive Pin	The value Yes indicates that the system is available exclusively for specific applications and spaces. The value No indicates that the system is available as a shared entity for all applications and spaces.
Pinned Applications	Displays the number of applications that are set to run on the host system.
Pinned Spaces	Displays the number of spaces mapped to the host system.

Pinned Applications section: This section displays an application and the corresponding space where it is deployed.

Pinned Spaces section: This section displays a space and the corresponding organization to which it belongs.

12.2.3.1.6 Maintaining Database Instances in XS Advanced

Applications deployed within a space in XS Advanced must be able to persist data. The XS Advanced Cockpit includes an option to maintain a database to persist application data. You can perform the following tasks to persist data:

- Create a database in XS Advanced
- Search for a database by name (or part of a name) in XS Advanced
- Display the status of all databases currently available in XS Advanced
- Enable a database for use with XS Advanced
- Disable a database in XS Advanced
- Delete a database in XS Advanced
- Map a database to an organization or space

Related Information

[Maintaining Tenant Databases in XS Advanced \[page 1713\]](#)

12.2.3.1.6.1 Create a Database

Context

Create a database to store data from applications available in XS Advanced.

Procedure

1. In the home navigation pane, choose *Tenant Databases*.
2. On the *Tenant Databases* screen, choose *New Tenant Database*.

The new database is created with the status *Creating*. Eventually, the status automatically changes to *Running* or *Not Running*.

Related Information

[Database Details \[page 1773\]](#)

12.2.3.1.6.1.1 Database Details

You provide the following details while creating a new logical database:

User Interface Element	Description
Name	Name of the database

User Interface Element	Description
Internal Port	<p>This is an optional field. You can provide details of a specific port on the system to run the tenant database. The port must adhere to one of the following rules:</p> <ul style="list-style-type: none"> • Accepts a value range and meets the condition ($\langle \text{port number} \rangle - (30000 + 40) \% 3 == 0$). For example, the range can vary from 30040 to 30100. <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin: 10px 0;"> <p>i Note</p> <p>The port value ($3 \langle \text{instance number} \rangle 00$) varies with the instance number. If the instance number is 08, the port value will be 30800. Confirm the instance number on your system and provide the respective port details.</p> </div> <ul style="list-style-type: none"> • Accepts the port number of the default index server.
Host	This is an optional field. If you want to host the tenant database on a specific system, provide the host details.
Tenant Database Password	Provide the password of the SYSTEM user to access the tenant database.
System Database User	Provide the SYSTEM user name as available on the SAP HANA system database.
System Database Password	Provide the password of the SYSTEM user of the SAP HANA system database.

12.2.3.1.6.2 Manage a Database

Context

To use a database once you have created it, you need to enable and map the database to an organization or space.

Procedure

1. In the home navigation pane, choose *Tenant Database*.

2. On the *Tenant Database* screen, you can perform any of the following tasks:

Task	Choose	Description
Make a logical database available for use	Enable	To use the database once you have created it, you need to enable it. The enabled database is available for use in XS Advanced. This is a toggle option. After enabling, the toggle option appears as <i>Disable</i> .
Use an organization or space with a database	 (Map)	To persist application data, you must map an organization or a space to the database. You can map an organization to a single database whereas a space can be mapped to multiple databases or vice versa.
Prevent a database from being used	 (Disable)	To ensure that a database is no longer available for use with XS Advanced, you must disable the database. The selected database is then no longer available for use.
Delete a database	 (Delete)	You can remove a database by deleting it.
Search for a specific database from a list of available databases	Search	You can use <i>Search</i> to find an existing database by name (or any part of the name). If you enter only part of a name, the search option filters the list of logical databases and displays only those whose names include the string you typed. For example, if you type DB in the search field, the list of databases displayed is restricted to any names that contain "DB"; for example, "MyTenantDB", "DB1", or "MyLogicalDB5".

12.2.3.17 Managing Trust Certificates in XS Advanced

For secure connections between SAP HANA systems and XS Advanced applications, each system needs to maintain a trusted certificate. This page displays a list of certificates currently used to establish trusted

connections. It also displays whether the current certificate being used is still active, expired, or is due to expire. You can perform the following tasks:

- Create a certificate
- Set a certificate for the default domain
- Manage certificates

Related Information

[Maintaining Trust Certificates in XS Advanced \[page 1706\]](#)

12.2.3.1.7.1 Upload a Certificate

Context

This option allows you to upload an existing certificate.

Procedure

1. In the home navigation pane, choose **Security** > **Trust Certificates**.
2. Choose **New Trust Certificate**.
3. In the **New Trust Certificate** dialog, provide the following details:

User Interface Element	Description
Alias	Name for the new trusted certificate. This name appears in the <i>Alias</i> column of the <i>Trusted Certificates</i> list.
Choose Certificate	Browse to the file containing the certificate.
Use with SAP HANA Service Broker	This option means that the certificate can be used when establishing the database connection.

12.2.3.1.7.2 Set Default Certificate

Context

You also have the option to set a certificate for the default domain.

Procedure

1. In the home navigation pane, choose ► *Security* ► *Trust Certificates* ►.
2. Choose *Set Platform Default Certificate*.
3. In the *Set Platform Default Certificate* dialog, provide the following details:

User Interface Element	Description
Choose Certificate	Browse to the file containing the certificate.
Choose Key	Private key in PKCS8 PEM format.

12.2.3.1.7.3 Manage Certificates

Context

You can view created certificates or delete any unwanted certificates.

Procedure

1. In the home navigation pane, choose ► *Security* ► *Trust Certificates* ►.
2. On the *Trust Certificates* screen, perform any of the following tasks:

Task	Choose	Description
View a certificate	 (View)	You can use this option to view the details of a certificate. This option is available underneath the Actions column.
Delete a certificate	 (Delete)	You can use this option to delete any unwanted certificates. This option is available under the Actions column.
Search for a specific certificate	Search	If the Trust Certificates screen lists a large number of certificates, you can use the search option to list a specific certificate. You can search for a certificate using any of the following certificate details: alias, usage, issuer, or valid details.

12.2.3.1.8 Maintaining Security in XS Advanced

During application development, developers create authorization information for business users. This information is made available to the administrators who complete the authorization setup. They are responsible for assigning authorizations to business users.

Developers store authorization information as design-time role templates in the security descriptor file [xs-security.json](#). Using the xsuaa service broker, they deploy the security information to a dedicated XS Advanced application. The XS Advanced administrators view the authorization information in role templates, which they use as part of the runtime configuration. The administrators use the role templates to build roles, which are aggregated in role collections. The role collections are then assigned to business users.

The tasks required to set up authorization artifacts in SAP HANA XS Advanced are performed using two distinct user roles: the application developer and the SAP HANA XS Advanced administrator. After the authorization artifacts have been deployed as role templates, the administrator of the SAP HANA XS Advanced application uses the role templates provided by the developers to build role collections and assign them to business users using the SAP HANA XS Advanced Cockpit.

i Note

To test authorization artifacts after deployment, developers can use the role templates to build role collections and assign authorizations to business users in the SAP HANA XS Advanced Cockpit.

The table below lists the sequence of tasks along with the user roles needed to set up authorization artifacts:

i Note

Application developers must have the *Space Developer* role in the required space. XS Advanced administrators must have the *XS_AUTHORIZATION_ADMIN* role.

Step	Task	User Role	Option to be Used
1	Specify the security descriptor file containing the functional authorization scopes for your application.	Application Developer	Text Editor
2	Create role templates for the XS Advanced application using the security descriptor file.	Application Developer	Text Editor
3	Create a service instance from the xsuaa service in XS Advanced using the service broker.	Application Developer	XS Advanced CLI tool / XS Advanced Cockpit
4	Bind the service instance to the XS Advanced application by including it in the manifest file.	Application Developer	Text Editor
5	Deploy the XS Advanced application.	Application Developer	XS Advanced CLI tool
6	If required, create a new role using role templates available within the application in the XS Advanced Cockpit.	XS Advanced administrator	Access applications within the space.
7	Create a role collection and assign roles to it.	XS Advanced administrator	Access applications within the space.
8	Assign the role collection to an SAML 2.0 identity provider or to SAP HANA database users.	XS Advanced administrator	Access applications within the space. Access SAML Identity Provider.
9	Assign the users to roles using the role collections	XS Advanced administrator	User interface of SAP HANA XS Advanced Cockpit

Related Information

[Control Access to Applications \[page 1769\]](#)

[Managing SAML Identity Providers in XS Advanced \[page 1782\]](#)

[Manage Users \[page 1756\]](#)

[Map Role Collections to SAML IDP \[page 1786\]](#)

12.2.3.1.8.1 Access to XS Advanced Cockpit

Context

The XS Advanced Cockpit displays all the options required to maintain the XS Advanced model runtime configurations. The options listed in the following table are available only to users who have been assigned the suitable role collection. The table shows the role collection required to use an option.

Options	Role Collection	Comments
Monitor	XS_CONTROLLER_ADMIN	Based on the role collection assigned, the user is permitted to perform some, most, or all operations in the application. The different role collections are: <ul style="list-style-type: none">• Admin : No access restrictions• User: Modify access within the assigned organization or space• Auditor: Read-only access within the assigned organization or space
Organization and Space Management	XS_CONTROLLER_USER or	
SAP HANA Service Broker Configuration	XS_CONTROLLER_AUDITOR	
Application Roles	XS_AUTHORIZATION_ADMIN or	Based on the role collection assigned, the user is permitted to perform some (or all) operations in the XS Advanced Cockpit. The different role collections are: <ul style="list-style-type: none">• Admin: Full admin edit access to the tool• Display: Read-only access to the tool
Trust Configuration	XS_AUTHORIZATION_DISPLAY	
Member Management	XS_USER_ADMIN	Based on the role collection assigned, the user is permitted to perform all operations in the application.

Options	Role Collection	Comments
SAP HANA Database Setup	XS_CONTROLLER_ADMIN	The user is permitted to perform all operations in the application.
Trusted Certificates		

As an administrator, you can see the default role collections available for your current role.

Procedure

1. In the home navigation pane, choose **Security > Role Collections**.
2. On the *Role Collections* screen, perform any of the following tasks:

Tasks	Choose	Description
Create a new role collection	New Role Collection	You can use this option to create a role collection.
Edit a role collection	 (Edit)	You can use this option to update descriptions of the role collections that you have created. You cannot edit a default role collection.
Delete a role collection	 (Delete)	You can use this option to delete a role collection that you have created. You cannot delete the default role collection.

12.2.3.1.8.2 Assign Roles to Role Collection

Context

As an administrator, you can control access to applications by providing role-based access to business users. To provide role-based access, you select the role templates and roles available for an XS Advanced application. You then assign the roles to role collections. The role collections are then assigned to business users to authorize them to access the application.

Procedure

1. In the home navigation pane, choose ► *Security* ► *Role Collections* ►.
2. On the *Role Collections* screen, choose the name of the role collection that you created.
3. On the *Overview* screen, choose *Add Role*.
4. In the *Add Role* dialog, select an application from the *Application Identifier*.
5. Choose the respective *Role Template* and *Role*.

Related Information

[Manage Users \[page 1756\]](#)

[Map Role Collections to SAML IDP \[page 1786\]](#)

12.2.3.1.9 Managing SAML Identity Providers in XS Advanced

You can configure an SAP HANA system to act as a service provider for XS Advanced applications that use single sign-on (SSO) authentication based on Security Assertion Markup Language (SAML) certificates.

The XS Advanced Cockpit includes Trust Configuration, which you can use to configure SAML identity providers at runtime. You must perform this step if you want your SAP HANA XS Advanced applications to use SAML assertions as the logon authentication method.

You can perform the following tasks:

- Create an SAML identity provider (IDP)
- Manage trust configurations
- Map role collections to an SAML IDP

12.2.3.1.9.1 Create an SAML Identity Provider

Prerequisites

You have configured the SAML identity provider with the assertion attribute **Groups** to be sent at runtime. Ensure that you have entered **G** in uppercase.

Context

You need to create an SAML Identity Provider (IDP) for an SAML service provider to authenticate users signing in to the application by means of single sign-on (SSO). SAP HANA supports the use of SSO authentication based on Security Assertion Markup Language (SAML) certificates.

Procedure

1. In the home navigation pane, choose **Security** > **Trust Configuration**.
2. On the **Trust Configuration** screen, choose **New Trust Configuration**.
3. Provide the required values.

Related Information

[Trust Configuration Details \[page 1783\]](#)

12.2.3.1.9.1.1 Trust Configuration Details

You can choose the **Edit** option to view details of an SAML IDP. To provide the required details for trust configuration, choose **New Trust Configuration**.

User Interface Element	Description
Upload	Choose this button to upload the XML format of the SAML certificate.
Origin Key	Unique name for the SAML identity provider. This value is mandatory.
Name	Name of the remote SAML identity provider.
Description	More details describing the SAML identity provider.
Status	State of the SAML identity provider.
Show SAML login link on login page	If the value is set as Yes , the SAML link appears on the logon page.
Link Text	The text representing the SAML logon URL on the logon page.

User Interface Element	Description
Metadata	The text containing the SAML certificate.
Show Details	If the contents of the XML document are valid, the parsing process extracts the information that needs to be inserted into the Origin Key, Subject, Entity ID, and Issuer fields, and the URL fields such as Single Sign-On URLs and Single Log-out URLs.
Parse	Choose this option to validate the SAML certificate if it is copied in XML format.
SingleSignOn URL (RedirectBinding)	URL of the IDP endpoint for SSO requests using SAML redirect binding.
SingleSignOn URL (PostBinding)	URL of the IDP endpoint for SSO requests using SAML post binding.
SingleLogout URL (RedirectBinding)	URL of the IDP endpoint for single logout (SLO) requests using SAML redirect binding.
SingleLogout URL (PostBinding)	URL of the IDP endpoint for single logout (SLO) requests using SAML post binding.

12.2.3.1.9.2 Manage Trust Configuration

Prerequisites

You have configured the SAML identity provider with the assertion attribute **Groups** to be sent at runtime. Ensure that you have entered **G** in uppercase. For information about how to configure the identity provider of your choice, see the topic [Federation Attribute Settings of Any Identity Provider](#) in the SAP Cloud Platform guide.

Context

You can modify the status or metadata details of SAML certificates. You can also delete unwanted certificates.

Procedure

1. In the home navigation pane, choose [Security](#) > [Trust Configuration](#).
2. On the [Trust Configuration](#) screen, perform any of the following tasks:

Task	Choose	Description
Update the status of an SAML certificate	 (Edit)	You can use this option to change the status of a certificate. This option is available underneath the Actions column.
Update the SAML metadata	Name of the SAML IDP > Edit	You can use this option to modify the SAML metadata.
Delete an unwanted SAML certificate	 (Delete)	You can use this option to delete an SAML certificate. This option is available underneath the Actions column.

12.2.3.1.9.3 Generate Metadata to Configure Identity Provider

Context

You need to generate a metadata (XML) file that contains information about the service provider (SAP HANA system) to configure the SAML identity provider. You do not have to be logged on to the SAP HANA XS Advanced Cockpit to generate the metadata.

Procedure

1. Access the XS Advanced Cockpit logon screen.
2. In the address bar, replace `/login` with `/saml/metadata`.
`<authorizationEndpoint>/saml/metadata`

Authorization end-point can be found by executing the command `xs -v` on the command line and looking for the key **authorizationEndpoint**. The metadata is downloaded. You can use the downloaded metadata to configure the SAML identity provider. The final step for providing SAML users with access to applications in the XS Advanced Cockpit is to map SAML assertion attributes with XSA attributes.

Related Information

[Map Role Collections to SAML IDP \[page 1786\]](#)

[The XS Command-Line Interface \[page 1653\]](#)

12.2.3.1.9.4 Map Role Collections to SAML IDP

Prerequisites

You have configured the SAML identity provider with the assertion attribute **Groups** to be sent at runtime. Ensure that you have entered **G** in uppercase.

Context

Users maintained in an SAML identity provider need authorization scopes to access XS Advanced applications. These scopes are contained in roles grouped in role collections. To provide authorization scope to SAML users logging on to an application using single sign-on (SSO), or using SAML-based SSO, you need to map SAML assertion attributes to XSA attributes.

Procedure

1. In the home navigation pane, choose **Security > Trust Configuration**.
2. On the *Trust Configuration* page, choose the name of the SAML IDP.
The *Overview* page of the SAML IDP appears.
3. In the trust configuration navigation pane, choose *Role Collection Mappings*.
4. Choose *New Role Collection Mapping*.
5. In the *Create Role Collection Mapping* dialog, provide the following details:

User Interface Element	Description
Role Collection	The dropdown list contains names of assertion-based role collections associated with the selected application.

User Interface Element	Description
Attribute	Displays the attribute defined in the selected application's security configuration (xs-security.json) file. Currently, the only attribute allowed is <i>Groups</i> .
Operator	The operator to be used along with the attribute in the specified rule. Currently, the only operator allowed is <i>equals</i> .
Value	The value of the attribute to be used for the rule that triggers the assignment of the selected role collection.

12.2.3.1.10 Managing Services in XS Advanced

Every space contains a service marketplace.

The marketplace displays SAP-approved services available for the space. Based on your role, you can access and use the services with your application. To use a service with your application, create an instance of the service. Once the instance is created, you can bind the instance to your application. You can also bind the application later using the *Service Instance* page, which lists instances of services.

User-provided services enable you to use services with your application that are not listed in the service marketplace, or custom services.

Related Information

[Maintaining Services in XS Advanced \[page 1692\]](#)

12.2.3.1.10.1 Create Service Instances

Context

To use a service from a service marketplace with an XS Advanced application, you need to create an instance of the service. Once the instance is created, you bind it to the application.

Procedure

1. Navigate to the Organization > Space in which the application is deployed.
2. Choose the application.

The application overview page appears.

3. In the navigation pane, choose *Service Bindings*.
4. In the service bindings pane, choose *Bind Service*.
5. Choose one of the following options:
 - *Service from the Catalog*: Use this option to bind to a service instance from the service marketplace.
 - *User-Provided Services*: Use this option to bind to a user-provided service instance.
6. Choose *Next*.
7. Select a service.
8. Choose *Next*.
9. Choose one of the following options:
 - *Create new instance*: This option creates a new instance of the service. Accordingly, choose a service plan.
 - *Re-use existing instance*: This option enables you to bind the application to an existing instance of the service.
10. Choose *Next*.
11. (Optional) Provide parameters that need to be passed to the application during binding, or provide a file with JSON format.
12. Choose *Next*.
13. Enter a name for the service instance.
14. Choose *Finish*.

The service instance is created and then bound to the application.

12.2.3.1.10.2 Manage User-Provided Services

Context

To use services that are not listed in the service marketplace with your application, you need user-provided services. User-provided service instances deliver service credentials to an application. A service instance for user-provided services behaves like the service instance created from the service marketplace.

Procedure

1. In the space navigation pane, choose **Services** > **User-Provided Services**.
2. In the **User-Provided Services** pane, choose **New Instance**.
3. Provide an instance name.
4. Enter the service credentials required to deliver to the application.
5. Choose **Save**.

The service instance is created. Similarly to services from the service marketplace, you need to bind the user-provided service instance to an application.

Related Information

[Create Service Instances \[page 1787\]](#)

12.2.3.1.11 Viewing Audit Logs in XS Advanced

You can use audit logs to check the various system logs (data access logs, security event logs, configuration change logs, or data modification logs) for the Java runtime environment using different criteria. You can use any of the logs for troubleshooting and identifying the cause of an issue.

Note

The XS Advanced Audit Log UI is no longer supported with SAP HANA 2.0, because the audit logs are now stored in the central SAP HANA audit log. To view logs in HANA 2.0, use SAP HANA tools to view the log entries.

Before you can use the XS Advanced Audit Log UI, you must make the following configuration settings:

- Configure your application for the audit log service.
- Set the necessary authorizations for your user.

To access the Audit Log tool, you can use either the XSA_ADMIN user (who already has the required scopes) or a different user, to whom you must assign a role collection that includes the AuditLogViewer role. You can choose the application and configure a role collection with the following information:

- Application name: auditlog-ui
- Application role: AuditLogViewer
- Template name: AuditLogViewer

Choose **User Management** to assign the custom role collection to the user.

Related Information

[Logging and Auditing in XS Advanced \[page 1836\]](#)

12.2.3.1.11.1 Start Audit Log Viewer from Cockpit

Context

Procedure

1. In the home navigation pane, choose *More*.
The *View Audit Logs* tile appears.
2. Choose the tile to display the *Audit Logs* interface.

12.2.3.2 Scheduling Jobs in XS Advanced

The Job Scheduler service enables you to create and schedule long-running operations or jobs.

In the SAP HANA XS advanced model, the Job Scheduler is an application service. The Job Scheduler service enables you to create and schedule long-running operations or jobs. This service is deployed during the installation of the SAP HANA XS advanced model.

i Note

For the correct functioning of Job Scheduler, you must deploy it in the UTC timezone.

The following table lists the sequence of tasks required to use an instance of the Job Scheduler service:

i Note

To configure and setup Job Scheduler you require specific roles and permissions.

Step	Task	Role
1	Configure the Service Broker for Job Scheduler	Space Developer
2	Create a Job Scheduler Service Instance	Space Developer

Step	Task	Role
3	Bind an Application to the Job Scheduler Service	Space Developer
4	Maintain jobs and job schedules	Administrator. To access and use the Job Scheduler Dashboard without having an administrator role, refer <i>The Job Scheduler Dashboard</i> topic.

Job Schedule Execution Mode

Job Scheduler supports the following **modes** for applications to execute a job:

- Synchronous Mode
Suitable for jobs that run for a short span of time, for example, an OData service end point
- Asynchronous Mode
Suitable for jobs that run for a long span of time, for example, end points which trigger batch processing

Job Schedule Execution Type

Job Scheduler provides the following **types** of schedules for a job:

- Recurring Schedule
Runs periodically at a specified time, dates, or interval. Recurring schedules can be created in the following ways:
 - The `repeatInterval` parameter:
Defines the interval in human-readable text (for example, "2 minutes"), which can be used to set up a recurring schedule. The repeat interval defines the gap between each run of the schedule.
 - The `cron` parameter:
Defines a `cron` expression (for example, "`cron`": "`* * * * *`") used to represent a set of times, when the job is executed.
 - The `repeatAt` parameter:
Defines the exact time, every day, when the job is executed.
- One-Time Schedule
Runs only once at the specified time. One-time schedules can be created in the following ways:
 - Human-readable text string:
A human-readable text string that defines the specific time for schedule execution (for example: "10 hours from now", "3.30pm", or "Friday at 2am")
 - Using a `Date` object, with a pre-defined format, for example,
`"startTime": {"date": "2015-10-20 4:30 +0000", "format": "YYYY-MM-DD HH:mm Z" }`
The string is checked against both IETF-compliant RFC 2822 timestamps and ISO-8601

Job Scheduler Access

Job Scheduler can be accessed and used in the following ways during application development:

- APIs:
The Job Scheduler service offers RESTful and client specific APIs for Java and Node.js. The administrator **scope** is required to use the Job Scheduler API to maintain run time configurations for jobs and job schedules.
- User Interface:
The *Job Scheduler Dashboard* is the tool used to manage the jobs and job schedules. Administrator authorization is required to maintain jobs and job schedules in the *Job Scheduler Dashboard*. For more information about permissions required to access the dashboard, see *The Job Scheduler Dashboard* topic.

i Note

You can program actions in any programming language or platform. The runtime also supports jobs created in the SAP HANA XS classic version.

Related Information

[Maintain Jobs and Job Schedules in XS Advanced \[page 1792\]](#)

[Job Scheduler REST API for XS Advanced \[page 1795\]](#)

[The Job Scheduler Dashboard \[page 1812\]](#)

12.2.3.2.1 Maintain Jobs and Job Schedules in XS Advanced

Maintain run time configurations for jobs and job schedules in SAP HANA XS advanced.

Prerequisites

- The service broker and the service instance for the Job Scheduler service are available.
- The application using the Job Schedule is deployed in the space and bound to the Job Scheduler service instance.
- You have the authorization scope for `POST`, `PUT`, and `DELETE` requests (for example, *jobscheduler.Admin*).
- To access the *Job Scheduler Dashboard*, you must have the authorization scopes defined in the roles grouped together in one of the following role collections:
 - `XS_CONTROLLER_ADMIN`
Full access: no access restrictions
 - `XS_CONTROLLER_USER`
Modify and read-only access
 - `XS_CONTROLLER_AUDITOR`

Read-only access

→ Tip

Role collections can be assigned to an SAP HANA user in SAP HANA studio by means of user parameters, for example, XS_RC_XS_CONTROLLER_ADMIN or XS_RC_XS_CONTROLLER_USER, or XS_RC_XS_CONTROLLER_AUDITOR.

Context

To maintain jobs and job schedules, you use the Job Scheduler REST APIs (for example, *Job Creation*, *Job Configuration*, or *Job Deletion*) as illustrated in the following examples.

i Note

The code examples are not always complete; they are intended for illustration purposes only.

Procedure

1. Create a new job.

Use the *Job Creation* API (POST /scheduler/Jobs), as illustrated in the example request:

```
POST /scheduler/jobs HTTP/1.1
Host: localhost:4242
Authorization: Basic YWJjOmRlZg==
Content-Type: application/json
Cache-Control: no-cache
{"name":"validateSalesOrder", "description": "cron job that validates sales
order requests", "action":"http://salesOrderApp.hana.acme.com:40023/
salesOrders/validate","active": true, "httpMethod":"PUT", "schedules":
[{"cron":"* * * * * */10", "description": "this schedule runs every 10
seconds", "data":{"salesOrderId":"1234"}, "active": true, "startTime":
{"date": "2015-10-20 04:30 +0000", "format": "YYYY-MM-DD HH:mm Z"}}]}
```

The response to the job-creation request should look like the following example:

```
{"name": "validateSalesOrder", "action":"http://salesOrderApp.hana.acme.com:
40023/salesOrders/
validate","active":true,"httpMethod":"PUT","description":"cron job that
validates sales order
requests","startTime":null,"endTime":null,"signatureVersion":0,"schedules":
[{"active":true,"startTime":"2015-10-20
04:30:00","endTime":null,"description":"every 10 seconds, every 2
minutes","data":{"salesOrderId":"1234"},"cron":"* * * * * */
10","type":"recurring","scheduleId":"cb5c9def-
e2a0-4294-8a51-61e4db373f99"}], "_id":3}
Headers:
Connection → keep-alive
Content-Length → 468
Content-Type → application/json; charset=utf-8
Date → Mon, 09 Nov 2016 09:08:53 GMT
ETag → W/"1d4-P7BnAm3yordzbrYyJtpalg"
```

```
Location → /scheduler/jobs/3
X-Powered-By → Express
```

2. Modify (configure) a new job.

Use the *Job Configuration* API (PUT /scheduler/Jobs), as illustrated in the example request:

```
PUT /scheduler/jobs/3 HTTP/1.1
Host: localhost:4242
Authorization: Basic YWJjOmRlZg==
Content-Type: application/json
Cache-Control: no-cache
{"active": true, "user": "abc", "password": "def", "httpMethod": "GET"}
```

The response to the job-configuration request should look like the following example:

```
{"success": true}
Headers:
Connection → keep-alive
Content-Length → 16
Content-Type → application/json; charset=utf-8
Date → Mon, 09 Nov 2016 09:30:36 GMT
ETag → W/"10-c2PoX+nt7m8FOksxlYjAhg"
X-Powered-By → Express
```

3. Delete an existing job.

Use the *Job Deletion* API (DELETE /scheduler/Jobs), as illustrated in the example request:

```
DELETE /scheduler/jobs/4 HTTP/1.1
Host: localhost:4242
Authorization: Basic YWJjOmRlZg==
Content-Type: application/json
Cache-Control: no-cache
```

The response to the job-deletion request should look like the following example:

```
{"success": true}
Headers:
Connection → keep-alive
Content-Length → 16
Content-Type → application/json; charset=utf-8
Date → Mon, 09 Nov 2016 09:30:36 GMT
ETag → W/"10-c2PoX+nt7m8FOksxlYjAhg"
X-Powered-By → Express
```

4. Create a new job schedule.

Use the *Job Schedule Creation* API (POST /scheduler/jobs/3/schedules), as illustrated in the example request:

```
POST /scheduler/jobs/3/schedules HTTP/1.1
Host: localhost:4242
Authorization: Basic YWJjOmRlZg==
Content-Type: application/json
Cache-Control: no-cache
{"repeatEvery": "2 hours", "data": {"order_id": "abcd"}, "active": true,
"description": "New Schedule", "startTime": {"date": "2016-04-21", "format":
"YYYY-MM-DD"}}
```

The response to the job-schedule creation request should look like the following example:

```
"repeatInterval": "2
hours", "repeatAt": null, "time": null, "cron": null, "data": "{ \"order_id\": \"abcd
\\\" }\", \"description\": \"New
Schedule\", \"type\": \"recurring\", \"active\": true, \"startTime\": \"2016-04-21
18:30:00\", \"endTime\": null, \"jobId\": 3, \"scheduleId\": \"0e29c67c-563e-4931-
af08-43acb10813e8\"}
Headers:
Connection → keep-alive
Content-Length → 274
Content-Type → application/json; charset=utf-8
Date → Mon, 09 Nov 2016 09:42:13 GMT
ETag → W/\"112-rdQSXHBVY0u6JNI/Wf0I7w\"
Location → /scheduler/jobs/3/schedules/0e29c67c-563e-4931-af08-43acb10813e8
X-Powered-By → Express
```

5. Delete an existing job schedule.

Use the *Job Schedule Deletion* API (DELETE /scheduler/jobs/3/schedules), as illustrated in the example request:

```
DELETE /scheduler/jobs/4 HTTP/1.1
Host: localhost:4242
Authorization: Basic YWJjOmRlZg==
Content-Type: application/json
Cache-Control: no-cache
```

The response to the job-schedule deletion request should look like the following example:

```
{ \"success\": true }
Headers:
Connection → keep-alive
Content-Length → 16
Content-Type → application/json; charset=utf-8
Date → Mon, 09 Nov 2016 09:51:39 GMT
ETag → W/\"10-c2PoX+nt7m8FOksxlyjAhg\"
X-Powered-By → Express
```

Related Information

[Job Scheduler REST API for XS Advanced \[page 1795\]](#)

[The Job Scheduler Dashboard \[page 1812\]](#)

[Scheduling Jobs in XS Advanced \[page 1790\]](#)

12.2.3.2.1.1 Job Scheduler REST API for XS Advanced

The Job Scheduler APIs enable applications to use the functionality provided in Job Scheduler.

The Job Scheduler-as-a-Service is a microservice component, which enables you to create, schedule, and run application tasks. The component exposes REST endpoints for interaction, with JSON as the format for data communication. The Job Scheduler API for SAP HANA XS advanced includes the commands listed in the following table. For more information about the configuration parameters required for the request, see the API documentation provided with the *Job Scheduler Dashboard* tool.

i Note

Access to the APIs is controlled by authorization scopes, for example, `admin` for `POST` and `PUT` requests, or `view` for `GET` requests. Scopes are built into roles, which can be assigned to users in role collections. The Job Scheduler REST APIs are protected with basic authentication.

An application, which has been bound to the Job Scheduler service and wants to interact with the Job Scheduler service, must extract the authentication credentials from the `<VCAP_SERVICES>` environment variable and use these credentials to call the REST APIs. To invoke the API, the user-authentication credentials must be encoded and passed in the "Authorization" header. If the credentials are not passed or they are passed wrongly, the APIs return a response with the status code "401- Unauthorized".

In this section, you can find information about the following topics:

- [Command Overview](#)
- [Human-Readable Dates](#)
- [Time Formats](#)

Command Overview

XS Advanced Job Scheduler REST API

API	Description	Required Scope
Job Creation	Used to create a job. Job creation can accept a collection of job schedules to be created.	<code>admin</code>
Job Configuration	Configure a job with updated run time information. The API can also be used to create a job if a Job with the Job Name in the URI segment, is not found.	<code>admin</code>
Job Deletion	Delete a job and purge all its run time information such as job schedules and logs.	<code>admin</code>
Job Schedule Creation	Create a job schedule for a specified job. All job configuration values (<code>Action URL</code> , <code>HTTP Method</code> , <code>User</code> , <code>Password</code> & <code>Job Activation Status</code>) are valid for the newly created schedule. A job schedule will only run if both the job and the schedule are active.	<code>admin</code>
Job Schedule Modification	Configure the run time information of a job schedule for a specified job. All job configuration values (for example: <code>Action URL</code> , <code>HTTP Method</code> , <code>User</code> , <code>Password</code> , and <code>Job Activation Status</code>) remain valid for the modified schedule.	<code>admin</code>
Job Schedule Deletion	Delete and purge run time information of the job schedule of the specified job. All related information like job schedule configurations and logs are purged. The processing of the schedule is also immediately stopped.	<code>admin</code>

API	Description	Required Scope
Bulk Job Schedule Activation	This is a utility API used to activate or deactivate all existing schedules of a job. This API triggers the immediate processing (or a halt in processing) of all job schedules for the specified job.	admin
Bulk Job Schedule Deactivation		
Job Details	Retrieve the saved details and configurations of a specified job. If the <code>displaySchedules</code> parameter is not provided, the schedules for the job are not returned and only the job details are returned.	view
Job Schedule Details	Retrieve the saved details and configurations of a specified job schedule & optionally the generated logs for the schedule.	view
Bulk Job Schedule Deletion	Delete and purge run time information of all the currently configured job schedules of the specified job. All related information like job schedule configurations and logs are purged. The processing of the schedules is also immediately stopped.	admin
Job Run Log Update	Used by applications, to inform the Job Scheduler about the status of an asynchronous, long-running job run.	admin

Job Creation

To create a job schedule, at least one of the fields `repeatAt`, `repeatEvery`, `cron` and `time` must be used. The response from the job creation API is a JSON body with the job details, including the ID of the job.

- **Route**

POST `/scheduler/jobs`

- **Response**

A JSON body containing the job details, including the ID of the job with status code "201-CREATED", if the call was successful. A location header with the relative path to the job-details is included in the response.

Sample Code

```
{
  "name": "validateSalesOrder",
  "description": "cron job that validates sales order requests",
  "action": "http://salesOrderApp.hana.ondemand.com:40023/salesOrders/validate",
  "active": true,
  "httpMethod": "PUT",
  "schedules": [
    {
      "cron": "* * * * */10 0",
      "description": "this schedule runs every 10 minutes",
      "data": {
        "salesOrderId": "1234"
      },
      "active": true,
      "startTime": {
        "date": "2015-10-20 04:30 +0000",
        "format": "YYYY-MM-DD HH:mm Z"
      }
    }
  ]
}
```

```

    }
  ]
}

Response:
{
  "name": "validateSalesOrder",
  "action": "http://<application-url>/action",
  "active": true,
  "httpMethod": "PUT",
  "description": "cron job that validates sales order requests",
  "startTime": null,
  "endTime": null,
  "signatureVersion": 0,
  "schedules": [
    {
      "active": true,
      "startTime": "2015-10-20 04:30:00",
      "endTime": null,
      "description": "every 10 seconds, every 2 minutes",
      "data": "{\"salesOrderId\":\"1234\"}",
      "cron": "* * * * * */10",
      "type": "recurring",
      "scheduleId": "schedule ID details"
    }
  ],
  "_id": 3
}

```

The job schedule creation request is defined with the parameters listed in the following table:

i Note

Parameters marked with an asterisk (*) are mandatory.

Job Creation: Request Body Fields

Request Field	Type	Description
name *	String	The unique name of the job to be created
		i Note If a job with the same name for the technical user credentials already exists, the job creation request fails.
description	String	Describes the user-defined job
action *	String	The fully qualified URL endpoint to be called when the job runs, for example: http://host.acme.com/app/call
active	Boolean	Defines if the job should be activated on creation. Allowed values are: <ul style="list-style-type: none"> • false (default) The job is in inactive mode on creation • true The job is activated on creation

Request Field	Type	Description
httpMethod	String	The HTTP method to be used to call the end-point URL for the job action . Allowed values are: GET, POST (default), PUT, and DELETE
startTime	Object	The start time for the job. If the start time is specified for the job, the scheduler checks if a start time is provided for the schedule as well. If a start time is provided for the schedule, it is used for determining the start of the schedule run. If no job-schedule start time is defined, the start time for the job is used. The date and time-formats must be specified as strings.
endTime	Object	The end time for the job. If the end time is specified for a job, the scheduler checks if an end time is provided for the schedule as well. If an end time is provided for the schedule, it is used for determining the end of the schedule run. If not, the end time for the job is used. The date and time-formats must be specified as strings.
schedules *	Array	The array of job schedule objects, to be created on job creation.

The `schedules` parameter can be used to provide details of the job schedule (as properties of each job schedule object); the following table lists the permitted properties:

Schedule Parameter Fields

Schedule Field	Type	Description
data	object	Optional data to be passed to the job action endpoint when invoked. Typically, the custom data is sent based on the HTTP method configured for invoking the end point URL, for example: <code>{"dataParam": "somevalue"}</code>
time	string or object	For one-time schedules, the parameter denoting the time at which the task executes. A human-readable text can be used to specify the time, for example, "3.30pm" or "tomorrow at 2am". If an object is used, the date and time-formats must be specified as strings.
repeatInterval	string	For recurring schedules, the parameter denoting the intervals when the schedule should run. The parameter supports the use of human readable formats.
repeatAt	string	For recurring schedules, the parameter denoting the exact time when the job schedule must run. A human-readable text can be used to denote a specified time, for example, "3.30pm" or "tomorrow at 2am", if the schedule runs repeatedly.
cron	string	For recurring schedules, the parameter denoting the cron pattern. It must be a valid cron tab format, for example: <code>"* * * * * */10"</code>
startTime	object	The time when the job scheduling should start. The date and time-formats must be specified as strings.

Schedule Field	Type	Description
endTime	object	The time when the job scheduling should end. The date and time-formats must be specified as strings
description	string	The user-provided description of the job schedule

Job Configuration

Configure a job with updated run time information. The API can also be used to create a job if a Job with the Job Name in the URI segment, is not found. If the API is being used to create a job, the parameters must conform to the same constraints as provided in the Job Creation API

- **Route**

```
PUT /scheduler/jobs/:jobId
```

```
PUT /scheduler/jobs/:jobName
```

“:jobId” is the ID of the job previously created using the Job Creation API. If the job name is used in the URI, it is first checked if the job with the name, exists. If no such named job exists, the API tries to create the job. If it does exist, the API configures the job with the details provided in the request body.

i Note

If the API is used to create a job, care must be taken to ensure that the job name in the request URI matches the name of the job in the request body. If the names do not match, an error is returned.

- **Response**

If the API finds an existing job, the response has a status code of “200-OK”, if the call was successful. The response has a status code of “201-CREATED”, if the API is used to create a new job; for new jobs, a location header containing the relative path to the job-details is returned in the response.

Sample Code

```
PUT /schedule/jobs/5 HTTP/1.1
content-type:application/json;charset=utf-8
host:https://scheduler.service.acme.com
content-length: 500
{"active": true, "user":"abc", "password":"def", "httpMethod": "GET"}
```

```
Response:
status: 200 OK
content-type: application/json; charset=utf-8
{"success": true}
```

Sample Code

```
PUT /schedule/jobs/jobwhichdoesnotexist HTTP/1.1
content-type:application/json;charset=utf-8
host:https://scheduler.service.acme.com
content-length: 500
{"name":"jobwhichdoesnotexist", "jobDescription": "greet the world
periodically", "action":"http://httpbin.org/basic-auth/abc/
```

```
def", "active": true, "httpMethod": "GET", "schedules": [{"repeatEvery": "2
minutes", "scheduleDescription": "every 2 minutes, run this schedule", "data":
{"time": "abc"}, "active": true}, {"cron": "* * * * *", "scheduleDescription":
"every 4 minutes, run this schedule", "data": {"time": "abc"}, "active":
false}]}
```

```
Response:
status: 201 CREATED
content-type: application/json; charset=utf-8
{"_id": 120, "name": "jobwhichdoesnotexist", "description": "", "action": "http://
httpbin.org/basic-auth/abc/
def", "active": true, "user": null, "httpMethod": "GET", "schedules":
[{"scheduleId": "b373469c-c6d4-4d5f-a002-c56f18455dc5", "description": "Default
Schedule", "data":
{"time": "abc", "type": "recurring", "active": true, "startTime": null, "endTime": nul
l, "repeatInterval": "2 minutes"}, {"scheduleId": "2f98471c-26de-4293-ae53-
e4a16e1513f5", "description": "Default Schedule", "data":
{"time": "abc", "type": "recurring", "active": false, "startTime": null, "endTime": nu
ll, "cron": "* * * * *"}]}
```

The job schedule configuration request is defined with the parameters listed in the following table:

Job Creation: Request Body Fields

Request Field	Type	Description
active	Boolean	Defines if the job should be activated on configuration. Allowed values are: <ul style="list-style-type: none"> false (default) The job is in inactive mode when configured true The job is active when configured
user	String	The name of the user account to run the configured job
password	String	The password for the user account to run the configured job
httpMethod	String	The HTTP method to be used to call the end-point URL for the job action . Allowed values are: GET, POST (default), PUT, and DELETE
startTime	Object	The start time for the job. If the start time is specified for the job, the scheduler checks if a start time is provided for the schedule as well. If a start time is provided for the schedule, it is used for determining the start of the schedule run. If no job-schedule start time is defined, the start time for the job is used. The date and time-formats must be specified as strings.
endTime	Object	The end time for the job. If the end time is specified for a job, the scheduler checks if an end time is provided for the schedule as well. If an end time is provided for the schedule, it is used for determining the end of the schedule run. If not, the end time for the job is used. The date and time-formats must be specified as strings.

Job Deletion

Delete a job and purge all its run time information such as job schedules and logs.

- **Route**
DELETE /scheduler/jobs/:jobId
- **Response**
If the call is successful, the response has a status code "200-OK" and includes a JSON response {"success": true}.

Sample Code

```
DELETE /schedule/jobs/:jobId HTTP/1.1
content-type:application/json;charset=utf-8
host:https://scheduler.service.acme.com
```

```
Response: Status: 200 OK
Content-Type: application/json;charset=utf-8
{"success":true}
```

Job Schedule Creation

Create a job schedule for a specified job. All job configuration values (Action URL, HTTP Method, User, Password & Job Activation Status) are valid for the newly created schedule. A job schedule will only run if both the job and the schedule are active.

- **Route**
POST /scheduler/jobs/:jobId/schedules
- **Response**
If the call is successful, the response has a status code of "201-CREATED". A location header with the relative path to the schedule-details, is returned in the response.

Sample Code

```
PUT /scheduler/jobs/3/schedules
{
  "repeatInterval":"2 hours",
  "active":true,
  "description":"New Schedule",
  "startTime": {
    "date": "2017-08-21",
    "format": "YYYY-MM-DD"
  }
}
```

```
Response:
{
  "repeatInterval": "2 hours",
  "repeatAt": null,
  "time": null,
  "cron": null,
  "data": "{\\"order_id\\":\\"abcd\\"}",
  "description": "New Schedule",
```

```

    "type": "recurring",
    "active": true,
    "startTime": "2015-04-20 18:30:00",
    "endTime": null,
    "jobId": 3,
    "scheduleId": "<schedule ID details>"
  }

```

Job Schedule Creation Parameters

Request Field	Type	Description
time	string or object	For one-time schedules, the parameter denoting the time at which the task executes. A human-readable text can be used to specify the time, for example, "3.30pm" or "tomorrow at 2am". If an object is used, the date and time-formats must be specified as strings.
repeatInterval	string	For recurring schedules, the parameter denoting the interval when the schedule should run. The parameter supports the use of human readable formats.
repeatAt	string	For recurring schedules, the parameter denoting the exact time when the job schedule must run. A human-readable text can be used to denote a specified time, for example, "3.30pm" or "tomorrow at 2am", if the schedule runs repeatedly.
cron	string	For recurring schedules, the parameter denoting the cron pattern. It must be a valid crontab format, for example: "* * * * * */10"
data	object	The parameter denoting optional data to be passed to the job action endpoint when invoked. Typically, the custom data is sent based on the HTTP method configured for invoking the end point URL, for example: {"dataParam": "somevalue" }
startTime	object	The time when the job scheduling should start. The date and time-formats must be specified as strings.
endTime	object	The time when the job scheduling should end. The date and time-formats must be specified as strings
active	Boolean	Defines if the job should be activated on configuration. Allowed values are: <ul style="list-style-type: none"> • false (default) The job is in inactive mode when configured • true The job is active when configured
description	string	The user-provided description of the job schedule

Job Schedule Modification

Configure the run time information of a job schedule for a specified job. All job configuration values (for example: Action URL, HTTP Method, User, Password, and Job Activation Status) remain valid for the modified schedule.

- **Route**
PUT /scheduler/jobs/:jobId/schedules/:scheduleId
- **Response**
If the call is successful, the response has a status code of 200– OK.

Calling this API stops further scheduling of the previously configured job schedule and, if activated, the processing for the newly configured schedule is started. This API cannot be used to change the scheduling mode for the job schedule. For example, if the schedule was created as a recurring “cron”-type schedule, it cannot be changed to a “repeatEvery”-type schedule. However, existing schedule values can be changed.

Sample Code

```
PUT /schedule/jobs/:jobId/schedules/:scheduleId HTTP/1.1
content-type:application/json;charset=utf-8
host:https://scheduler.service.acme.com
content-length: 500
{"description": "Edited Schedule", "startTime": {"date": "2013-02-08
09:30:26.123"}, "endTime": {"date": "2015-06-08 09:30:26.123"}, "active":
true, "cron": "* * * * *"} }
```

```
Response:
Status: 200 OK
Content-Type: application/json; charset=utf-8
{"scheduleId":"80e23846-734e-4b4b-a130-159a492ec482","name":"greet the
world3","data":{"time":"abc"},"type":"recurring","priority":
0,"action":"http://httpbin.org/basic-auth/abc/
def","nextRunAt":"2015-04-23T03:58:21.358Z","startTime":"2013-02-08T04:00:26.1
23Z","endTime":"2015-06-08T04:00:26.123Z","active":true,"description":"Edited
Schedule","jobId":"136","cron":"* * * * *"} }
```

Job Schedule Modification Parameters

Request Field	Type	Description
time	string or object	For one-time schedules, the parameter denoting the time at which the task executes. A human-readable text can be used to specify the time, for example, “3.30pm” or “tomorrow at 2am”. If an object is used, the date and time-formats must be specified as strings.
repeatInterval	string	For recurring schedules, the parameter denoting the interval when the schedule should run. The parameter supports the use of human readable formats.
repeatAt	string	For recurring schedules, the parameter denoting the exact time when the job schedule must run. A human-readable text can be used to denote a specified time, for example, “3.30pm” or “tomorrow at 2am”, if the schedule runs repeatedly.

Request Field	Type	Description
<code>cron</code>	string	For recurring schedules, the parameter denoting the <code>cron</code> pattern. It must be a valid <code>crontab</code> format, for example: <code>"* * * * * */10"</code>
<code>data</code>	object	The parameter denoting optional data to be passed to the job action endpoint when invoked. Typically, the custom data is sent based on the HTTP method configured for invoking the end point URL,for example: <code>{"dataParam": "somevalue"}</code>
<code>startTime</code>	object	The time when the job scheduling should start. The date and time-formats must be specified as strings.
<code>endTime</code>	object	The time when the job scheduling should end. The date and time-formats must be specified as strings
<code>active</code>	Boolean	Defines if the job should be activated on configuration. Allowed values are: <ul style="list-style-type: none"> <code>false</code> (default) The job is in inactive mode when configured <code>true</code> The job is active when configured
<code>description</code>	string	The user-provided description of the job schedule

Job Schedule Deletion

Delete and purge run time information of the job schedule of the specified job. All related information like job schedule configurations and logs are purged. The processing of the schedule is also immediately stopped.

⚠ Caution

This API removes all the run time configuration information of the job schedule, irrespective of whether the schedule is active or not.

- **Route**

```
DELETE /scheduler/jobs/:jobId/schedules/:scheduleId
```

- **Response**

If the call is successful, the response has a status code, "200-OK" and includes a JSON response `{"success": true}`.

📄 Sample Code

```
DELETE /scheduler/jobs/:jobId/schedules/:scheduleId HTTP/1.1
content-type:application/json;charset=utf-8
host:https://scheduler.service.acme.com
```

```
Response: {"success": true}
Status Code: 200 OK
```

Bulk Job Schedule Activation/Deactivation

This is a utility API used to activate or deactivate all existing schedules of a job.

- **Route**
POST /scheduler/jobs/:jobId/schedules/activationStatus
- **Response**
If the call is successful, the response has a status code, "200-OK" and includes a JSON response {"success": true}.

Sample Code

```
POST /scheduler/jobs/:jobId/schedules/activationStatus HTTP/1.1
content-type:application/json;charset=utf-8
host:https://scheduler.service.acme.com
{"activationStatus": true}
```

```
Response: {"success": true}
Status Code: 200 OK
```

Bulk Job Schedule Activation Parameters

Request Field	Type	Description
activationStatus	Boolean	The desired activation status of the job schedules for the job. Allowed values for the activation status are: <ul style="list-style-type: none">• false (default) All job schedules for the specified job should be deactivated• true All job schedules for the specified job should be activated

Job Details

Retrieve the saved details and configurations of a specified job.

- **Route**
GET /scheduler/jobs/:jobId?displaySchedules=true
GET /scheduler/jobs?jobId=:jobId&displaySchedules=true Route
GET /scheduler/jobs?name=:jobName&displaySchedules=true
- **Response**
If the call is successful, the response has a status code, "200-OK" and includes a JSON response with the schedule details, for example: {"schedules": [{"data": {"time": "abc"}, "type": "recurring", "repeatInterval": "2 minutes", "active": false, "startTime": null, "endTime": null, "repeatAt": null, [...]}.

Sample Code

```
GET /scheduler/jobs/:jobId?displaySchedules=true HTTP/1.1
content-type:application/json;charset=utf-8
```

```
host:https://scheduler.service.acme.com
```

```
Response:
Status: 200 OK
Content-type: application/json;charset=utf-8
{"schedules":[{"data":{"time":"abc"},"type":"recurring","repeatInterval":"2
minutes","active":false,"startTime":null,"endTime":null,"repeatAt":null,"sched
uleId":"0d3b4cc1-0f7b-4ee6-ab12-63d474b900f2","description":"Default
Schedule"}, {"data":{"time":"abc"},"type":"recurring","cron":"* * * *
*","active":false,"startTime":null,"endTime":null,"repeatAt":null,"scheduleId"
:"1b1bb70f-cada-46c9-9974-a7a1b87ba24f","description":"Default
Schedule"}],"name":"greet the world2","description":"","action":"http://
httpbin.org/basic-auth/abc/
def","user":null,"httpMethod":"GET","active":false,"_id":111}
```

Job Details Parameters

Request Field	Type	Description
displaySchedules	Boolean	Display details of the job schedules for the job. Allowed values for the job details are: <ul style="list-style-type: none">• false Do not display details of job schedules for the specified job• true Display details of job schedules for the specified job
jobId	String	The job ID needed to query for the job details. This can be passed as a URI segment parameter or as a query parameter.
name	String	The job name needed to query the job details. This can be passed as a query parameter

Job Schedule Details

Retrieve the saved details and configurations of a specified job schedule & optionally the generated logs for the schedule. Either `:jobId` or `:name` is required to invoke this API. If `displayLogs` is not provided, the logs for the schedule are not returned and only the schedule details are returned.

- **Route**

```
GET /scheduler/jobs/:jobId/schedules/:scheduleId?displayLogs=true
```

- **Response**

If the call is successful, the response has a status code, "200-OK" and includes a JSON response with the schedule details, for example: `{"data":`

```
 {"time":"abc"},"type":"recurring","repeatInterval":"2
minutes","plannedTime":"2015-04-19T15:12:44.000Z","active":true,"startTime":null
,"endTime":null,"nextRunAt": "2017-08-11 10:00:00","repeatAt":null, [...]}.

```

Sample Code

```
GET /scheduler/jobs/112/schedules/550d1b96-8002-4d0d-850e-368aaa591671?
displayLogs=true
HTTP/1.1 content-type:application/json;charset=utf-8
host:https://scheduler.service.acme.com
```

```

Response:
Status: 200 OK
Content-Type: application/json; charset=utf-8
{"data":{"time":"abc"},"type":"recurring","repeatInterval":"2
minutes","plannedTime":"2015-04-19T15:12:44.000Z","active":true,"startTime":nu
ll,"endTime":null,"nextRunAt": "2017-08-11 10:00:00","repeatAt":null,"logs":
[{"text":null,"httpStatus":null,"executionTime":null,"status":"SCHEDULED","sch
eduleTime":"2015-04-19T15:10:53.000Z","completionTime":null}], "scheduleId":"55
0d1b96-8002-4d0d-850e-368aaa591671","description":"Default Schedule"}

```

Job Schedule Details Parameters

Request Parameter	Type	Description
displayLogs	Boolean	Controls whether the API should return (<code>true</code>) all the generated logs for the job schedule or not (<code>false</code>)

Bulk Job Schedule Deletion

Delete and purge run time information of all the currently configured job schedules of the specified job. All related information like job schedule configurations and logs are purged. The processing of the schedules is also immediately stopped.

⚠ Caution

This API removes all the run time configuration information of the job schedule, irrespective of whether the schedule is active or not.

- Route**
`DELETE /scheduler/jobs/:jobId/schedules`
- Response**
 If the call is successful, the response has a status code, "200-OK" and includes a JSON response `{"success": true}`.

☰ Sample Code

```

DELETE /scheduler/jobs/:jobId/schedules HTTP/1.1
content-type:application/json;charset=utf-8
host:https://scheduler.service.acme.com

```

```

Response: {"success": true}
Status Code: 200 OK

```

Job Run Log Update

Inform the Job Scheduler about the status of an asynchronous, long-running job run. This API must be invoked by the application after the asynchronous execution of the job has completed, with the status of the job run and optionally some text about the job execution.

⚠ Caution

This API must be invoked by the application after the **asynchronous** execution of the job has completed, with the status of the job run and optionally some text about the job execution.

- **Route**

PUT /scheduler/jobs/:jobId/schedules/:scheduleId/runs/:runId

i Note

Parameters marked with an asterisk (*) are mandatory.

Job Run Log Update Parameters

Request Parameter	Type	Description
success *	Boolean	Indicates that the job run was successful (<code>true</code>) or failed (<code>false</code>)
message	String	Additional log/text about the job run

Human Readable Dates

The job scheduler for XS advanced supports human readable dates and ranges for the parameters `time`, `repeatAt` and `repeatEvery`, which are used for configuring job schedules. The job scheduler uses an embedded English language date parser for this facility. Valid human readable strings for the parameters are shown below:

i Note

The date parser expects a valid readable string; invalid strings will either throw parser errors or cause the job scheduling to happen inconsistently.

Date and Time Parameters

Parameter	Comments	Examples
<code>time</code>	Designates a particular timestamp for running a job schedule. If an invalid string is provided, the scheduler falls back to the current timestamp and runs the schedule immediately. The following example strings are valid for the <code>time</code> parameter:	"10 hours from now" "20 minutes from now" "in 2 hours" "tomorrow at 4pm" "next week monday at 5am" "9pm tonight" "3.30pm"

Parameter	Comments	Examples
<code>repeatAt</code>	Represents a convenient way to create daily timestamp-based schedules. The string should designate a particular timestamp for repeatedly running a job schedule. This follows the same pattern as the recommendations for the "time" parameter, barring a few discrepancies. While the text for the "time" parameter must denote something concrete and in the future, the 'repeatAt' must designate a timestamp, which is valid and constant daily. If an invalid string is used, the scheduler falls back to the current timestamp and runs the schedule immediately.	"4.40pm" "18.40" "6.20am" "17.20:30 " "09:30:26.123+07:00"
<div style="background-color: #f0f0f0; padding: 10px; border-left: 2px solid #0070c0;"> <p>i Note</p> <p>Second-based precision can sometimes be inaccurately timed; timezones must be specified using the offset (in hours), for example, "+07:00"</p> </div>		
<code>repeatInterval</code>	The string should designate a interval to repeat the job execution. Word strings for denoting the numeric value are not supported yet. For example, for "twenty minutes", use "20 minutes" to denote the interval. Supported time-units for this parameter are "years", "months", "weeks", "days", "hours", "minutes", "seconds".	"10 hours " "2 days " "3 seconds"

Date and Time Formats in Job Schedule Parameters

The date-time parameters for job schedules (for example, `startTime`, `endTime`, and `time`) can be passed as objects , with the mandatory `date` field denoting the date as a string and an optional `format` field denoting a date-time format for correctly parsing the user-provided date value. If the parameters are passed as strings, they must be valid date representations, in either the ISO-8601 or IETF-compliant RFC 2822 formats. For object representations, the following rules apply:

- Date field as input**
 If only the date field is provided as input, the string is checked against both IETF-compliant RFC 2822 time stamps and ISO-8601. If the date string is of an unknown format, the parser displays an error. For ISO-8601 compliant dates, calendar dates (for example, "2013-02-08"), week dates ("2013-W06-5"), ordinal dates ("2013-039") and time-based dates ("2013-02-08 09+07:00") are all supported.
- Date string format**
 If the format of the date string is customized, an optional format string can be passed. The allowed parsing tokens are as described in the following table:

Date and Time Parameters

Input Token	Example	Description
YYYY	2014	4 digit year

Input Token	Example	Description
YY	14	2 digit year
Q	1-4	Quarter of year. Sets month to first month in quarter
M MM	1-12	Month number
MMM MMMM	January- Dec	Month name in locale
D DD 1- 31		Day of month
Do	1st- 31st	Day of month with ordinal
DDD DDDD	1-365	Day of year
X	1410715640.579	Unix Timestamp
x	1410715640579	Unix Timestamp (ms)
gggg	2015	Locale 4 digit week year
gg	15	Locale 2 digit week year
w ww	1- 53	Locale week of year
e	1-7	Locale day of week
GGGG	2015	ISO 4-digit week year
GG	15	ISO 2-digit week year
W WW	1- 53	ISO week of year
E	1-7	ISO day of week
H HH	0 -23	24 Hour Time
h hh	1-12	12 hour time used with 'a A'
a A	am pm	Post or ante meridiem
m mm	0 -59	Minutes
s ss	0 -59	Seconds
S	0 -9	Tenths of a second
SS	0 -99	Hundredths of a second
SSS	0 -999	Thousandths of a second
Z ZZ	+12:00	Offset from UTC as +-HH:mm, +-HHmm, or Z

Date-Time Format Examples

- `startTime`
`"startTime": {"date": "2015-10-20 4:30 +0000", "format": "YYYY-MM-DD HH:mm Z"}`
4.30 UTC on 20th Oct 2015
- `endTime`
`"endTime": {"date": "2015-W06-5"}`
Friday, February 06, 2015
- `time`
`"time": {"date": "2010-10-20 4:30", "format": "YYYY-MM-DD HH:mm"}`
4.30 Local Time (the timezone for the scheduler service is considered here)

12.2.3.2.1.2 The Job Scheduler Dashboard

The Job Scheduler dashboard enables you to manage job schedules for a service instance.

The dashboard lists the available jobs. Select a job to create a schedule or to view existing schedules.

How to access the Job Scheduler dashboard

1. Get the Dashboard URL

To access the Job Scheduler dashboard, you can connect remotely using secure shell (SSH) to the SAP XS Advanced server and perform the following steps:

1. List the applications running on the server using the command `xs apps`.
2. In the *URL* column, identify the relevant URL for the application ("jobscheduler-dashboard" for Job Scheduler).
3. Copy the URL to any Web browser to launch the application.

i Note

If a valid certificate is not available, the Web browser indicates an issue with the certificate. To resolve the issue, add the required certificate.

4. Enter the logon credentials to access the application.

2. Permission to Access the Dashboard

The administrator role with the `XS_CONTROLLER_ADMIN` role template contains the `controller.admin` permission. This role template gives you full access to perform all the administration-related tasks. If you do not have the administrator role, you need the roles and role collection templates listed in the table below to view or modify settings.

Task	Roles Required	Role Collection Template Required
Open the dashboard and only view the job listing page	Org Manager Space Developer	jobscheduler_viewer_template
<p>i Note You cannot create jobs.</p>		
Configure settings	Org Manager Space Developer	jobscheduler_config_template
Create or edit jobs or schedules	Org Manager Space Developer	jobscheduler_admin_template

i Note

If you already have the permission to create or edit jobs or schedules, you still need **jobscheduler_config_template** to configure settings.

The SAP HANA administrator creates a role collection and adds the Job Scheduler roles to the role collection. For more information, see *Maintaining the SAP HANA XS Advanced Model Run Time* in the *SAP HANA Administration Guide*.

Dashboard Screens

The following table contains the various screens available on the dashboard and their descriptions:

Screen	Description
Configuration	<p>Enables you to maintain the global configuration required for a specific Job Scheduler service instance.</p> <ul style="list-style-type: none"> • Max. Invocation Attempts: The number of attempts made by the Job Scheduler to reach the job action endpoint before it deactivates the job. If the Job Scheduler fails repeatedly to reach the endpoint, it sets the job to inactive. The default value of this parameter is 3. • Asynchronous Execution timeout (ms): The duration (in milliseconds) that the Job Scheduler waits for a response for the asynchronous job from the application endpoint. If the application does not provide a response in the specified duration, the run status is set to COMPLETED/UNKNOWN.

Screen	Description
Jobs	Lists all the jobs created for a specific service instance. You can delete a job or navigate to the job details by choosing the job name.
Overview	Displays the details of the selected job. You can edit a job.
Schedules	Enables you to create and configure schedules for a job. To access schedules, select a job listed on the dashboard. The Schedules screen is displayed. Select a schedule to see the history and logs corresponding to the schedule. To display the run logs of a schedule, choose Logs .
Action History	Maintains the history of a job or schedule for a specific job.

Related Information

[Scheduling Jobs in XS Advanced \[page 1790\]](#)

12.2.4 XS Advanced User Management

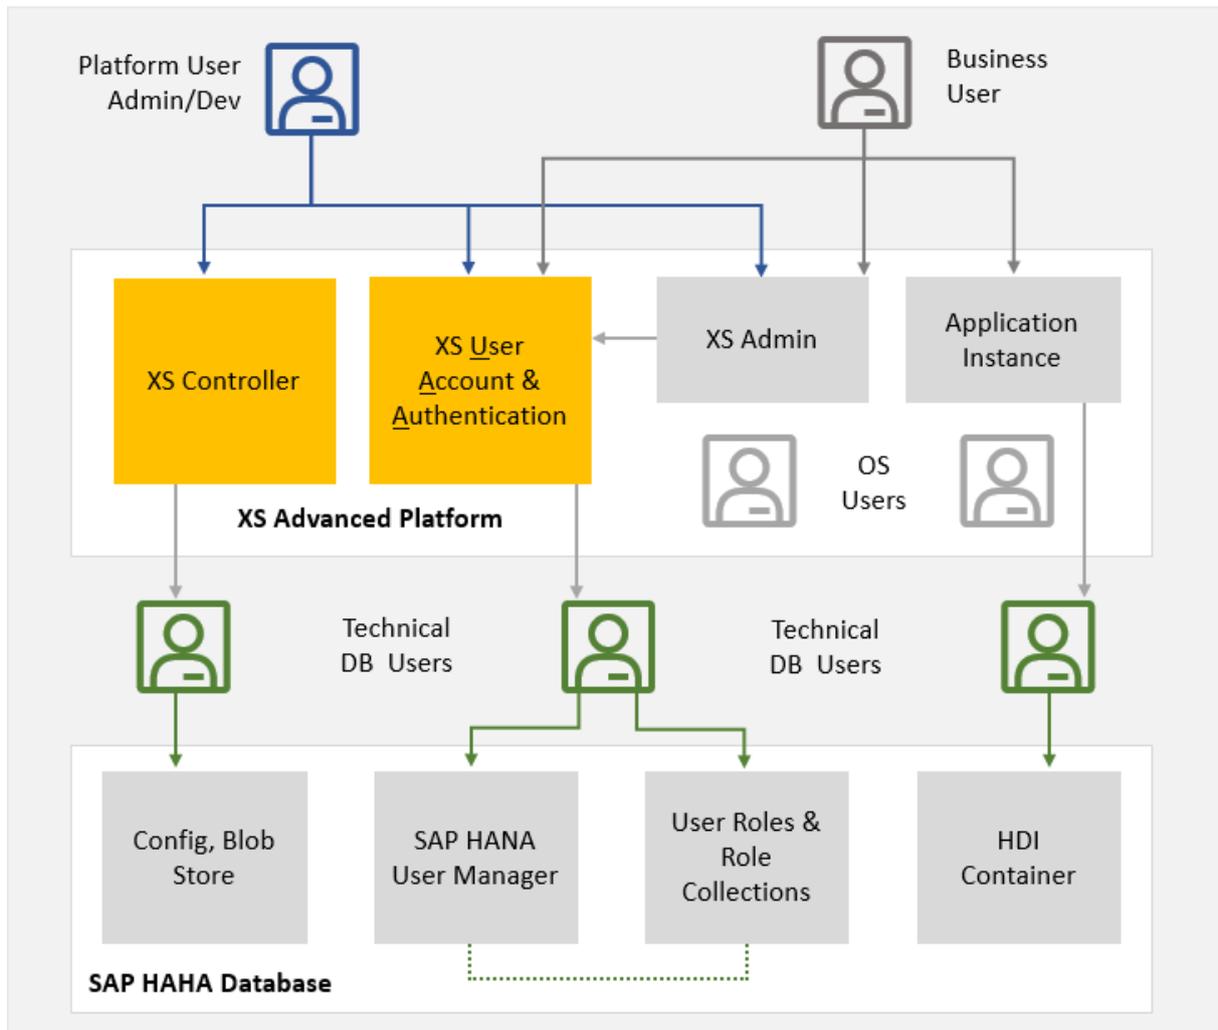
Manage XS advanced users with the tools integrated into the XS advanced platform.

SAP HANA XS, advanced model, requires the configuration and availability of a number of different users, for example: application developers, application users, and XS advanced administrators, as well as technical users, which includes technical database users: database users with restricted permissions. You can manage the users and their roles and scopes with a selection of tools, for example, the `xs` command-line interface, the [XS Admin Tools](#), or an Identity Provider (IdP).

XS Advanced User Overview

In traditional application servers, user information is kept in a local user store. However, although the SAP HANA XS advanced platform uses the underlying SAP HANA user store as an Identity Provider (IdP) by default, it is possible to integrate an external IdP such as SAP ID Service or SAP Cloud Identity. It is also possible to configure a custom IdP provided the IdP implements the SAML 2.0 standard.

The **User Account and Authentication (UAA)** service represents the central platform service for the management and authentication of users, as illustrated in the following diagram:



User Account and Authentication Service (UAA)

User information, such as first name, last name, user ID and user privileges, is provided in the form of signed OAuth2 access tokens the central UAA issues when a client logs in successfully. For more information about the authentication procedure, see the section on XS advanced user authentication.

XS Advanced User Categories

XS advanced users access the back-end instances typically through end-user interfaces such as a Web browser or command-line tools. Unlike technical users, application users and some types of system user can be also identified by personal data such as name, e-mail address, and so on. As the same identity provider is the basis for all of XS advanced users, an application user may also be granted developer privileges and the other way around.

XS advanced users who have their source in the SAP HANA user store (default) are typically restricted users with no access to SAP HANA database schemas. In contrast, applications and server components use

technical SAP HANA users with certain access privileges. The platform passes these credentials to applications, enabling them to execute SQL statements, if the XS advanced user has sufficient privileges. Decoupling XS advanced users from technical users is the precondition for leveraging external IdPs, even though XS advanced users are also SAP HANA users by default. As technical SAP HANA users are generated by the platform in the background, you typically won't use them to interact with the system.

The XS Advanced User Account and Authentication (UAA) service provides the authentication end point for individual users who need to interact either with SAP HANA XS, advanced model, or with applications deployed and running on the XS advanced model platform. Although such users are often referred to simply as **XS advanced users**, they can have the following roles and areas of responsibility:

- [Platform users \[page 1816\]](#)
- [Application or business users \[page 1816\]](#)
- [Operating-system users \[page 1816\]](#)
- [Technical database users \[page 1817\]](#)

XS Advanced Platform User

Platform users are administrators or developers who are assigned to one or more specific organizations or spaces in the XS advanced platform. An XS advanced administrator user (for example, `XSA_ADMIN`) is allowed to perform any platform operation in any organization or space. However, it is also possible to maintain additional platform user **roles** and use them to restrict the type of access granted to certain users for particular organizations or spaces. XS advanced “administration users” are system users who manage the configuration of the XS advanced application server components, and in particular the XS Controller.

XS advanced platform users are SAP HANA users who have been assigned to a specific XS advanced role collection. Non-administrator platform users can also be managed by means of an external Identity Provider (IdP).

→ Tip

You can use the `xs` command-line interface to maintain XS advanced platform users. For more information, see *Maintaining Platform Users in XS Advanced with the XS CLI* in *Related Information* below.

XS Advanced Application (Business) User

Often referred to as “business users”, application users interact with application instances deployed to and running on the XS advanced run-time platform. Application users are also referred to as business users, for example, employees, customers, and so on.

Application users can be identified by personal data such as name or e-mail address, and this data along with other credentials are stored in a user store, for example, an Identity Provider (IdP); any request to log on to an XS advanced application is managed by the XS advanced User Account and Authentication service (UAA). Authorization scopes (defined in user roles) are granted to application users to restrict or enable access to particular data.

XS Advanced Operating-System Users

In the context of XS advanced, the following predefined operating system users are available by default:

- `<sid>adm`
Operating-system and administrative SAP HANA system user who owns all platform services as well as the system's file storage.
- `sap<sid>xsa`

Operating-system user required for staging and running applications in the pre-configured `SAP` space.

- `<sid>xsa`

Operating-system user required for staging and running applications in the pre-configured `PROD` space.

The `<SID>adm` operating system user exists to provide an operating system context. From the operating system perspective, the operating system administrator is the user that owns all SAP HANA files and all related operating system processes. Certain administration operations require the operating system user's credentials, for example, starting or stopping the system.

i Note

XS advanced application files are also owned by the `*xsa` operating-system users `sap<sid>xsa` and `<sid>xsa`.

XS Advanced Technical Database User

A technical database user does not correspond to a real person and should be used for administrative tasks such as creating objects and granting privileges for a particular application. For example, an application server may log on to the SAP HANA database using a dedicated technical database user. In the context of XS advanced, technical SAP HANA users are generated by the platform in the background.

For XS advanced, the technical user `SYS_XS_RUNTIME` owns the XS Advanced Controller's SAP HANA schema, which contains the Blob Store, Config Store, and Secure Store. Similarly, the technical user `SYS_XS_UAA` owns the SAP HANA schema provided for the User Account and Authentication (UAA) for user management

Additional technical database users are created on demand and as required for application-specific purposes. For example, the `SBSS_*` users are created as a result of an application-service binding. XS advanced also makes use of a number of `USR_*` users, too; `USR_*` users are created by the SAP HANA Service Broker for the service plans `schema`, `securestore`, and `sbss`. Similar to the predefined users created when binding an application to an HDI container, `USR_*` users are used by applications to access their schema. For more information, see *Predefined Users* in *Related Information* below.

i Note

With HANA 2.0 SPS 03, the SAP HANA Service Broker no longer uses the `SBSS_*` prefix for HDI container users. Instead, these HDI container users have the name of the corresponding HDI container as the prefix. For example, for users created during service binding, the following format is used:
`<HDI_Container_Name>_<GUID>_DT` (design-time access) or `<HDI_Container_Name>_<GUID>_RT` (run-time access). Binding users are assigned the role `PUBLIC` by default.

Related Information

[Predefined Users in XS Advanced \[page 1818\]](#)

[Predefined XS Advanced Database Roles \[page 1822\]](#)

[Configure HDI Parameters \[page 1469\]](#)

[Maintaining Platform Users in XS Advanced \[page 1671\]](#)

[Maintaining Organizations and Spaces in XS Advanced \[page 1663\]](#)

[Maintaining the XS Advanced Run-time Environment with a Graphical User Interface \[page 1753\]](#)

12.2.4.1 Predefined Users in XS Advanced

The installation of the XS advanced application server creates a small set of predefined users that enable the operation of the underlying system.

The system's super user (<sid>adm) needs to be available in order to manage the life cycle of the system. Similarly, an administrative XS advanced system user (XSA_ADMIN by default) is necessary to perform the initial setup of the application server, for example, granting other users the privilege to create spaces in a dedicated organization and so on. Technical database users are created during installation for all server components that need to persist data in SAP HANA schemas.

This topic contains information about the following types of predefined users in XS advanced:

- [Predefined XS Advanced System Users](#)
- [Predefined SAP HANA Technical Users](#)
- [Predefined XS Advanced Operating-System Users](#)

Predefined XS Advanced System Users

The table below lists the predefined XS advanced system users that are necessary for operating the XS advanced application server. First, an administrative user named XSA_ADMIN is required for the XS advanced Controller; this administrative user configures the application server at a global level. Non-administrative users of the XS advanced Controller are not allowed to perform administration tasks, for example, uploading custom certificates, adding custom buildpacks, or registering platform service URLs. Bear in mind that, although the credentials for the technical users for the SAP HANA Service Broker and UAA Broker are generated automatically during installation, the XSA_ADMIN user is created interactively with a user-defined password. As a first-level administrator user with irrevocable privileges, the XSA_ADMIN has unlimited access to the XS advanced Controller and therefore needs to be handled carefully.

→ Recommendation

- Keep the number of people with XSA_ADMIN credentials as small as possible. Where possible, delegate specific tasks such as space management to users with less privileges instead.
- Avoid creating other powerful users with privileges similar to XSA_ADMIN.
- Change the XSA_ADMIN password at regular intervals and avoid sharing the same password.

User ID	User Type	Description
XSA_ADMIN	XS advanced user	Administrative user for the XS advanced application server with unlimited access to XS advanced Controller API
HDI_BROKER_CONTROLLER	Technical user	Technical user for the SAP HANA Service Broker API

Although the technical users in this table are created in the SAP HANA database, and database authentication checks are used to confirm the technical users' credentials, the technical users are not used to connect to the SAP HANA database.

Predefined Technical SAP HANA Users

Most of the server agents require a data store in the SAP HANA database and therefore need secure access to schemas. For this reason, a dedicated technical SAP HANA user is generated for each such schema, and the credentials of the technical SAP HANA user are passed to the server agent. As the management of technical users is performed at the infrastructure level, end users typically do not interact with these users. The technical users listed in the following table are used to connect to the SAP HANA database with a specific set of conditions.

User ID	Service	Type	Description
HDI_ADMIN_USER	SAP HANA Broker	Technical database user	Owns SAP HANA schema of SAP HANA Service Broker
HDI_BROKER_CONTROLLER	SAP HANA Broker	Technical database user	Has authorization to access the service broker API of SAP HANA broker
SYS_XS_HANA_BROKER	SAP HANA Broker	Technical database user	Owns the SAP HANA Service Broker's SAP HANA schema
SYS_XS_HANA_BROKER_INTERNAL	SAP HANA Broker	Technical database user	Has authorization to execute stored procedures for creating users, and so on.
SYS_XS_INSTANCE_MANAGER_ADMIN_USER	Instance Manager	Technical database user	Owns SAP HANA schema of the Instance Manager
SYS_XS_INSTANCE_MANAGER_BROKER_USER	Instance Manager	Technical database user	Has authorization to access service broker API of Instance Manager
SYS_XS_OID_USER	OIDC	Technical database user	Owns the SAP HANA schema for the OpenID Connect provider
SYS_XS_OID_USER_SEC	OIDC	Technical database user	Owns the SAP HANA secure store for the OpenID Connect provider
SYS_XS_RUNTIME	Controller	Technical database user	Owns the Controller's SAP HANA schema containing BlobStore, ConfigStore and SecureStore

User ID	Service	Type	Description
SYS_XS_SBSS	SAP HANA Broker	Technical database user	Owns SAP HANA schema containing procedures to generate user passwords in a secure manner; used by the SAP HANA Service Broker
SYS_XS_SYSTEMDB_INFO	Controller	Technical database user	Has authorization to access database system catalog and configuration
SYS_XS_UAA	UAA	Technical database user	Owns the UAA's SAP HANA schema for user management
SYS_XS_UAA_SEC	UAA	Technical database user	Owns the UAA's SAP HANA secure store for user credentials
SYS_XSA	Installer	Owns SAP HANA schema containing a unique tenant ID	Owns SAP HANA schema containing a unique tenant ID
_SYS_DI	HDI	Technical database user	Owns all HDI SQL-based APIs, for example all API procedures in the _SYS_DI schema and API procedures in containers
_SYS_DI_*_CATALOG	HDI	Technical database user	Technical users used by the HDI to access database system catalog tables and views
_SYS_DI_SU	HDI	Technical database user	Technical superuser of the HDI created at installation time
_SYS_DI_TO	HDI	Technical database user	Owns transaction and connections of all internal HDI transactions

Technical Users for HDI Schema-Based Containers

The deployment of database objects with SAP HANA Deployment Infrastructure (HDI) is based on a container model where each container corresponds roughly to a database schema. Each schema, and the database objects deployed into the schema, are owned by a dedicated technical database user.

For every container deployed, a new technical database user and schema with the same name as the container are created. Additional schemas and technical users required for metadata and deployment APIs are also created.

For example, for a container named *s*, HDI creates the following users:

- *s*:
The user who is the owner of the container schema *s*

- User `S#DI`:
The user who is the owner of the schema `S#DI` containing metadata and deployment APIs
- User `S#OO`:
The user who is the owner of database objects in schema `s`
- Users `_DI#S#METADATA_COM_SAP_HANA_DI_<metadata>`:
The users who are the owners of schemas containing build plug-in metadata

These technical users are used internally by HDI only. They are created as restricted database users who do not have any privileges by default (not even the role `PUBLIC`). They cannot be used to log on to the database.

For more information, see *Maintaining HDI Containers* in the *SAP HANA Developer Guide (For SAP HANA XS Advanced Model)*.

Technical Users for Default Application Services

XS advanced applications can make use of a number of services managed by a service broker. To make use of a service, an instance of the service must be created and the application must be bound to the specified service instance. Several services are available by default; they are installed with the XS advanced run-time platform.

The installation of the following default application services results in the creation of a number of internal technical users:

- `Product-Installer`
Used for the installation and installation management of applications
- `Deploy-Service`
Used in the technical deployment of applications packaged in multi-target application (MTA) archives

The operation of binding these services to an application generates a technical user and random password according to the following naming convention `USR_<generated_ID>`. These technical users are required to make database schemas available for applications. For every combination of application and schema, such a technical user is created.

In addition, the `Job-Scheduler` service, used to create and schedule long-running operations in the XS advanced environment, uses an HDI container with the `SBSS_` prefix and a randomly generated name. The above-mentioned HDI schemas and users will be created for this container.

For more information, see *The SAP HANA XS Advanced Services: Deployment Infrastructure* in the *SAP HANA Developer Guide (For SAP HANA XS Advanced Model)*.

Predefined XS Advanced Operating System Users

Ultimately all platform services are made up of operating-system (OS) artifacts such as OS processes, network sockets, and file storage. Since operating systems come with their own user management features, these artifacts are out of necessity owned by OS users. Consequently, the XS advanced application server cannot be run without at least one OS user, although dedicated XS advanced users are able to perform the majority of the operational tasks.

The installation procedure creates the “super” OS user `<sid>adm` for the entire SAP HANA system. As the owner of all operating-system processes, the `<sid>adm` user is very powerful from a security perspective. For this reason, we strongly recommend that you limit the number of people with `<sid>adm` credentials as far as possible.

The following XS advanced platform services launch new processes at run time:

- Execution Agents start application instances.
- The application “Stager” spawns processes running build packs during application staging.

In both cases, custom code comes to execution. If these processes ran as the system's `<sid>adm` user, the whole system could be compromised. To prevent this, the platform generally spawns external processes with OS users that are attached to the application's space. To support this approach, the initial setup includes OS user `<sid>xsa` user for the `PROD` space and OS user `sap<sid>xsa` for the `SAP` space.

The following table summarizes the operating-system users that are available immediately after installation:

User ID	Type	Description
<code><sid>adm</code>	OS user	Administrative SAP HANA system user who owns all platform services as well as the system's file storage
<code><sid>xsa</code>	OS user	OS user for staging and running applications in the pre-configured <code>PROD</code> space
<code>sap<sid>xsa</code>	OS user	OS user for staging and running applications in the pre-configured <code>SAP</code> space

Related Information

[Predefined XS Advanced Database Roles \[page 1822\]](#)

[XS Advanced User Management \[page 1814\]](#)

12.2.4.2 Predefined XS Advanced Database Roles

Several predefined database roles are necessary for the operation of the XS advanced model application server.

i Note

The following roles are SQL-based roles that are available in the catalog of the SAP HANA database.

Role	Description
_SYS_DI_OO_DEFAULTS	<p>This role contains the set of default privileges that are granted to all HDI container object owner users (<container>#OO users). SAP HANA DI uses this role internally to grant default privileges instead of using the PUBLIC role. It contains only privileges to SYS views where additional security checks apply.</p> <p>The role contains SELECT privileges on the views: SYS.DUMMY, SYS.PROCEDURES, SYS.PROCEDURE_PARAMETERS, SYS.TABLES, SYS.TABLE_COLUMNS.</p> <p>This role is not intended to be granted to database users.</p> <div data-bbox="804 741 1401 853" style="background-color: #f0f0f0; padding: 5px;"> <p>i Note Do not extend this role in a production system.</p> </div>
SYS_XB_SBSS_VIEWER	<p>This role contains selected privileges for monitoring the status of the Service Broker Security Support (SBSS) component.</p> <p>The SBSS component provides service brokers with functions for creating, validating, and deleting the credentials they need for service bindings. Credential handling is achieved by creating restricted database users with secure random passwords.</p> <p>Specifically, this role contains read access to the SBSS component version table, in addition to read access to the SBSS bindings table that lists the credential names that have already been created with the SBSS API as well as some meta data for the bound credentials.</p> <p>This role is intended only for support users so they can query information such as SBSS version, number of credentials, names of services brokers that called the SBSS API.</p> <div data-bbox="804 1480 1401 1592" style="background-color: #f0f0f0; padding: 5px;"> <p>! Restriction This role does not grant access to any SBSS credentials.</p> </div>

Related Information

[Predefined Users in XS Advanced \[page 1818\]](#)

[XS Advanced User Management \[page 1814\]](#)

12.2.5 XS Advanced System Configuration Parameters

The XS advanced platform can be configured with a selection of system parameters.

Most system parameters are set to a default value during the installation of the XS advanced platform component. If necessary, however, a small selection of XS advanced system parameters can be changed after the installation, too, for example, to meet the requirements of your system landscape and usage patterns. If you need to change the default SSL port or set up a fail-over router, you can change system properties with SAP Cockpit, SAP HANA Studio, or by editing the configuration file itself, for example, `xscontroller.ini`.

Note

The XS advanced run-time platform reads its configuration from the SAP HANA initialization (`.ini`) files. To apply any changes made to the XS advanced system parameters, restart the XS advanced run time.

This section contains information about the following XS advanced configuration parameters:

- [Global XS Advanced Configuration Parameters \[page 1824\]](#)
- [XS Advanced Controller Configuration \[page 1825\]](#)
- [XS Advanced Execution Agent Configuration \[page 1830\]](#)

Global XS Advanced Configuration Parameters

The following table lists the system-wide parameters that can be changed in the `global.ini` file and suggests possible scenarios where a change might be necessary.

XS Advanced Global System Parameters `global.ini`

Property Name	Ini-File Section	Default Value	Description	Reason for Change
<code>basepath_xsa_app-workspace</code>	<code>persistence</code>	<code>hana/shared/ <SID>/xs/ app_working</code>	The path to the working directory of XS advanced. The directory is used for the execution of staging processes and applications. All files in this directory exist only temporarily. If the working directory is changed, it is not necessary to copy the directory contents to the new working directory.	If the startup times for staging processes and applications are slow, change the path to a location on the local file system with a good I/O throughput rather than the shared file system.

Property Name	Ini-File Section	Default Value	Description	Reason for Change
xsa_sizing	system_information	M	The sizing profile of the XS advanced platform. For more information, see <i>Platform Sizing</i> in <i>Related Information</i> below.	You have a development system or you expect a lot of business users on the production system.

Configuration Parameters for the XS Advanced Controller

If you need to change the behavior of the XS advanced Controller, the following table lists the parameters that you can modify in the `xscontroller.ini` file and suggests possible scenarios where a change might be necessary.

XS Controller System Parameters (xscontroller.ini)

Property Name	Ini-File Section	Default Value	Description	Reason for Change
syslog	audit	false	Write the audit log to the file system (default) or into the system log.	You want the audit log messages written into the system log.
default_domain	communication	Set during installation	The domain used for the URLs of platform components, for example, the XS Controller or XS UAA. The domain is also used by default for all application URLs (routes).	Set up a fail-over router. Change the application URLs. Set up an additional reverse proxy in front of XS advanced.

Property Name	Ini-File Section	Default Value	Description	Reason for Change
router_port	communication	3<instance nr>33	<p>The port the XS advanced platform router is listening on if the <code>hostnames</code> routing mode is enabled. The port number is used for every URL or route.</p> <div data-bbox="927 629 1150 1249" style="border: 1px solid #ccc; padding: 5px;"> <p>i Note</p> <p>If you configure the router port to a port < 1024, you must set the appropriate file permissions for the <code>icmbnd</code> binary at <code>/hana/shared/<SID>/xs/router/webdispatcher</code>, for example, <code>chown root:sapsys</code> and <code>chmod 4705</code>.</p> </div>	<p>You want to use the default SSL port 443 to reach your applications.</p> <div data-bbox="1169 521 1394 723" style="border: 1px solid #ccc; padding: 5px;"> <p>i Note</p> <p>This requires some additional configuration steps.</p> </div>
router_portrange_start	communication	51000	<p>The range of ports used by the XS advanced platform to generate port-based routes if no specific port is defined.</p>	<p>You want to create more than 500 routes.</p> <p>You want to run several XS advanced systems on one host.</p> <div data-bbox="1169 1480 1394 1892" style="border: 1px solid #ccc; padding: 5px;"> <p>→ Tip</p> <p>The port-range modification should be completed before the installation of the second system. For more details, see SAP Notes 2507070 and 2243156.</p> </div>

Property Name	Ini-File Section	Default Value	Description	Reason for Change
router_portrange_end	communication	51500		A change to the port range after installation of XS advanced is not recommended because it will not have any effect on the current routes.
router_https	communication	true	Determines whether the Platform Router provides HTTPs or HTTP endpoints for URLs and routes	You fully trust the network from which applications are accessed and want to disable HTTPs due to the additional performance overhead of SSL.
internal_https	communication	true	Determines whether the XS advanced platform components communicate using HTTPs and client certificate authentication based on the SAP HANA System PKI. If the property is turned off, no authentication between the XS Controller and the Execution Agents is performed.	You fully trust the network from which applications are accessed and want to disable HTTPs due to the additional performance overhead of SSL.

Property Name	Ini-File Section	Default Value	Description	Reason for Change
jdbc_ssl	communication	false	<p>Determines whether the XS advanced platform components and applications should use SSL encrypted communication with the SAP HANA database.</p> <div data-bbox="927 663 1152 1167" style="background-color: #f0f0f0; padding: 5px;"> <p>i Note</p> <p>Additional configuration steps are required to ensure that all applications are capable of establishing an SSL connection successfully. For more information see <i>Maintaining Trust Certificates in Related Information</i> and SAP Note 2300943.</p> </div>	<p>You want to:</p> <ul style="list-style-type: none"> • Ensure connections with the HANA database are encrypted • Enable SSL enforcement for connections to the SAP HANA database.

Property Name	Ini-File Section	Default Value	Description	Reason for Change
jdbc_ssl_validate_certificate	communication	false	<p>Determines whether the XS advanced applications should validate the SAP HANA SSL certificate; this requires additional configuration, which is described in <i>Maintaining Trust Certificates</i>.</p> <div style="border: 1px solid orange; padding: 5px; margin-top: 10px;"> <p>! Restriction</p> <p>This property only effects Java-based applications running in the XS advanced run-time environment.</p> </div> <p>SSL validation cannot be switched off for other (non-Java) application run time environments. The XS advanced platform components always validate the SSL certificate, as they are using the System PKI to establish trusted SSL communication with the SAP HANA database.</p>	You want to ensure that the encrypted connection of applications with the SAP HANA database are not vulnerable to man-in-the-middle attacks.
jdbc_ssl_certificate_hostname	communication	-	<p>Specifies the host name to use when validating the SSL certificate during the SSL handshake.</p> <div style="border: 1px solid orange; padding: 5px; margin-top: 10px;"> <p>! Restriction</p> <p>This property only effects Java-based applications running in the XS advanced run-time environment.</p> </div>	The host name in the certificate does not match the host name the XS advanced run time environment uses to connect to the SAP HANA database.

Property Name	Ini-File Section	Default Value	Description	Reason for Change
mdc_dispatcher	general	true	Determines whether MDC Dispatcher or the sudo command is used to start application instances as different operating-system (OS) users.	You want to ensure the sudo-based user switch is used, regardless of the SAP HANA version.
<div style="background-color: #e0e0e0; padding: 5px; border: 1px solid #ccc;"> <p>→ Tip</p> <p>For more information see SAP Note 2243156.</p> </div>				
default_space_user	general	<sid>xsa	Determines the operating-system user defined as the default user for a newly created space. The specified default space user is used to start application instances in the target space.	You do not want to use the <sid>xsa user created by the command hdb1cm as the default space user.

Configuration Parameters for the XS Advanced Execution Agent

If you need to change the behavior of the XS advanced Execution Agent, the following table lists the parameters that you can modify in the `exeagent.ini` file and suggests possible scenarios where a change might be necessary.

XS Execution Agent System Parameters (exeagent.ini)

Property Name	Ini-File Section	Default Value	Description	Reason for Change
listen_portrange_start	communication	50000	The ports from this range are used to assign the internal port to an application instance. These ports are not supposed to be reached externally. Externally the application instances are reached via the Webdispatcher.	You want to run more than 1000 apps You want to run several XS advanced systems on one host.
<div style="background-color: #e0e0e0; padding: 5px; border: 1px solid #ccc;"> <p>→ Tip</p> <p>Since ports are dynamically assigned to applica-</p> </div>				

Property Name	Ini-File Section	Default Value	Description	Reason for Change
listen_portrange_end	communication	50999		tion instances during startup, you can change this property after installation. For more information see SAP Note 2507070 .
exclusive	communication	false	If you pin applications to specific Execution Agents, this property controls whether the Execution Agent can start only its pinned applications or other applications, too.	You want to prevent applications from running on the same host. → Tip For more information, see <i>Maintaining Host Pinning</i> in <i>Related Information</i> .

Related Information

[Maintaining Domains in XS Advanced \[page 1708\]](#)

[Maintaining Host Pinning \[page 1717\]](#)

[Maintaining Trust Certificates in XS Advanced \[page 1706\]](#)

[Platform Sizing in XS Advanced \[page 1839\]](#)

12.2.6 Backup and Recovery in XS Advanced

Back up and restore XS advanced persistent data.

Since XS advanced stores most of its persistent data within the SAP HANA database, you can use the SAP HANA tools and mechanisms to backup and restore XS advanced. There is, however, some persistent data, namely the file-system service, which is not automatically included in the SAP HANA backup. XS advanced provides tools to store this information within the database manually and restore it back to the file system.

! Restriction

Earlier versions of SAP HANA do not automatically back up the file-system secure store that is used by XS advanced. For detailed instructions about how to back up XS advanced using earlier SAP HANA versions, see SAP Note [2300937](#).

Performing a Backup of XS Advanced

The procedure for backing up XS advanced model differs according to the options you chose when you set up XS advanced during installation, for example, whether you installed XS advanced in the system database or a tenant database, as described in the following sections.

→ Tip

For more information about installing XS advanced in the System database or in a custom database, see *XS Advanced Database Setup Options* in *Related Information* below

Backup with XS Advanced Installed in the System Database

In this scenario, the XS advanced platform services (for example XS Controller) store all persistent data, such as application and service metadata, run times, build packs, and application droplets within the System database of the SAP HANA Platform. Applications running on XS advanced might store their persistent data within the System database or within a tenant database, depending on the respective configuration of the org and space the application is deployed to. For more information, see *Maintaining Tenant Databases* in *Related Information*.

If you want to backup a XS advanced system you must always backup the System database and all of the tenant databases, which are used within XS advanced. In case you want to be able to recover the backup to a fresh-installed system, it is recommended to backup all tenant databases (not just the ones used by XS advanced), to prevent errors with tenant databases, which are known to the System database, but do not exist in the system after recovery. For more information about how to back up the System database and tenant databases, see *SAP HANA Backup* in *Related Information* below.

⚠ Caution

During a backup operation inconsistencies can occasionally occur between the SAP HANA service-instance metadata stored in the System database by the SAP HANA Service Broker and the actual state of the SAP HANA service instance (for example, HDI containers, schemas, and users) within the tenant databases. It is not possible, however, for inconsistencies in the data stored within a SAP HANA service instance to occur.

For this reason, and to avoid backup inconsistencies, it is recommended to shut down XS advanced before starting the backup process. Shutting down XS advanced before the backup is not required if you are sure that no SAP HANA service instances are created, deleted, bound, or unbound in between the backup of the System database and the backup of the respective tenant databases. This is usually the case in productive systems, when no applications are updated or initially deployed during the backup. It is important to bear in mind that inconsistencies in the data stored within an SAP HANA service instance cannot occur.

1. Disable XS advanced services.

Shut down XS advanced by running the following command as `<sid>adm` user:

```
$ XSA disable
```

2. Back up the file-system service.

To include the File-System Service (FSS) in an XS advanced backup operation, run the `backup-fss` command as `<sid>adm`, as shown in the following example:

```
$ XSA backup-fss
```

i Note

Applications can store data in other backing services provided by service brokers, for example, the file-system service. The file-system service is used by SAP Web IDE to store some of its persistent data. XS advanced provides the tools needed to store the contents of the file-system service within the System database and restore the contents at a later point in time. It is recommended to back up the XS advanced file-system before you perform the backup of the System database.

3. Back up the databases.
First back up the System database and then all the tenant databases. For more details of how to back up the databases, see *SAP HANA Backup in Related Information*.
4. Enable XS advanced services again.
Restart XS advanced by running the following command as `<sid>adm` user:

```
$ XSA enable
```

Backup with XS Advanced Installed in a Tenant Database

In this scenario, the XS advanced platform services (for example XS Controller) store all persistent data, such as application and service metadata, run times, build packs, and application droplets within a tenant database of the SAP HANA Platform. Applications running on XS advanced might store their persistent data within the same or a different tenant database, depending on the respective configuration of the organization and space the application is deployed to. For more information, see *Maintaining Tenant Databases in Related Information*.

If you want to back up a XS advanced system you must always back up all of the tenant databases that are registered with XS advanced. For more information about how to backup the tenant databases, see *SAP HANA Backup in Related Information* below.

⚠ Caution

If XS-advanced-related data is spread across different tenant databases, the consistency of the backed up data cannot be ensured if XS advanced services are running during the backup. For this reason, it is recommended to shut down XS advanced during the back up operation.

It is not necessary to shut down XS advanced during the back up process if you are sure that no SAP HANA service instances are created, deleted, bound or unbound in between the backup of the system database and the backup of the respective tenant databases. This is usually the case in productive systems, when no applications are updated or initially deployed during the backup. It is important to bear in mind that inconsistencies in the data stored within an SAP HANA service instance cannot occur.

1. Disable XS advanced services.
If more than one tenant databases are registered with XS advanced, shut down XS advanced by running the following command as `<sid>adm` user:

```
$ XSA disable
```

i Note

If only one tenant database is registered with XS advanced, it is **not** necessary to shut down the XS advanced service. This tenant database contains XS-advanced-related data and can be backed up consistently without disabling XS advanced first.

2. Back up the file-system service.

To include the File-System Service (FSS) in an XS advanced backup operation, run the `backup-fss` command as `<sid>adm`, as shown in the following example:

```
$ XSA backup-fss
```

Note

Applications can store data in other backing services provided by service brokers, for example, the file-system service. The file-system service is used by SAP Web IDE to store some of its persistent data. XS advanced provides the tools needed to store the contents of the file-system service within the System database and restore the contents at a later point in time. It is recommended to back up the XS advanced file-system before you perform the backup of the System database.

3. Back up the tenant databases.

To back up the tenant databases registered with XS advanced, first display a list of all the registered tenant databases, for example, by logging into SAP HANA system as `<sid>adm` and running the following command:

```
$ XSA list-tenants
```

Backup all tenant databases displayed in the command output. For more information about how to back up SAP HANA databases, see *SAP HANA Backup* in *Related Information* below.

→ Tip

Make a note of the tenant database containing the XS advanced platform persistence printed by the `list-tenants` command; this is the database that needs to be selected after recovery.

```
DB name: MYTENANT  
[...]  
XS advanced platform persistence: YES  
[...]
```

4. Enable XS advanced services again.

If you disabled the XS advanced services before starting the backup operation, then after the backup is complete, restart XS advanced by running the following command as `<sid>adm` user:

```
$ XSA enable
```

Restoring XS Advanced from a Backup Image

The procedure for restoring XS advanced model differs according to the options you chose when you set up XS advanced during installation, for example, whether you installed XS advanced in the system database or a tenant database.

→ Tip

For more information about installing XS advanced in the System database or in a custom database, see *XS Advanced Database Setup Options* in *Related Information* below

Restoring XS Advanced Installed in the System Database

If you want to restore XS advanced that was installed in the system database, you need to perform the following high-level steps:

1. Disable XS advanced services.

To avoid problems during the recovery operation, it is necessary to shut down XS advanced on the system, where the XS advanced backup should be restored. To shut down XS advanced, run the following command as `<sid>adm` user:

```
$ XSA disable
```

2. Restore XS-advanced-related data.

To restore the XS advanced system from a backup copy, you need to restore the System database as well as **all** of the tenant databases that are used within XS advanced. The order in which you perform these steps is mandatory; restore the System database first and then all of the tenant databases. If you are restoring the backup to a freshly installed system, bear in mind that you need to restore all the tenant databases that are known to the System database. If you do not adhere to this rule, you might experience errors with tenant databases that are known to the System database, but do not exist in the system.

→ Tip

For more information about how to recover the System database and tenant databases see *SAP HANA Recovery in Related Information*.

3. Restore the file-system service.

After you have successfully recovered the System database and all tenant databases you can restore the file-system service to XS advanced, too, using the following command as `<sid>adm`:

```
$ XSA restore-fss
```

4. Enable XS advanced services again.

To enable XS advanced, run the following command as `<sid>adm` user:

```
$ XSA enable
```

Restoring XS Advanced Installed in a Tenant Database

If you want to restore XS advanced that was installed in a tenant database, you need to perform the following high-level steps:

1. Disable XS advanced services.

To avoid problems during the recovery operation, it is necessary to shut down XS advanced on the system, where the XS advanced backup should be restored. To shut down XS advanced, run the following command as the `<sid>adm` user:

```
$ XSA disable
```

2. Restore XS-advanced-related data.

To restore the XS advanced system from a backup copy, you need to restore **all** tenant databases that contain XS advanced data.

→ Tip

For more information about how to recover a tenant database, see *SAP HANA Recovery in Related Information*.

3. Point the XS advanced platform services to the tenant database containing the XS advanced platform data you just restored, for example, by running the following command as `<sid>adm` user

```
$ XSA select-xsa-runtime-db <tenant database name>
```

For details of how to identify the database containing the XS advanced platform data, see the command `XSA list-tenants` in *Backup with XS Advanced Installed in a Tenant Database* above.

4. Restore the file-system service.

After you have successfully pointed the XS advanced platform services to the correct tenant database, you can restore the file-system service to XS advanced, too, using the following command as the `<sid>adm` user:

```
$ XSA restore-fss
```

5. Enable XS advanced services again.

To restart all XS advanced services, run the following command as the `<sid>adm` user:

```
$ XSA enable
```

Related Information

[SAP Note 2300937](#) 

[SAP HANA Backup \[page 1245\]](#)

[SAP HANA Recovery \[page 1331\]](#)

[Maintaining Tenant Databases in XS Advanced \[page 1713\]](#)

12.2.7 Logging and Auditing in XS Advanced

Set up and use logs and traces in XS advanced.

The XS Advanced platform writes different types of traces several different files in the following host-specific SAP HANA trace directories:

```
/usr/sap/<SID>/HDB<SN>/<host>/trace
```

The XS advanced platform keeps access logs concerning requests sent to XS advanced applications via the platform router centrally in the following directory:

```
/hana/shared/<SID>/xs/controller_data/controller/router/webdispatcher/logs
```

XS Advanced Platform Tracing

Each XS Advanced platform component, for example, the XS Controller, the XS Execution Agent, or the XS User Account and Authentication service (UAA) writes detailed tracing to log files located in the SAP HANA trace directory of the host the service is running on. The following files contain the most recent log entries:

- `xscontroller_0.log`
- `xsexeeagent_0.log`
- `xsuaaserver_0.log`

i Note

Log files are closed and renamed (rotated) if they exceed a certain size.

Log-file rotation means that you can find older traces in files with a higher index incorporated into the file name, for example, `xscontroller_1.log` or `xscontroller_2.log`. The platform retains the five (5) most recent trace files for each XS advanced service. In addition, startup information is written to the following files in the SAP HANA trace directory:

- `xscontroller.out`
- `xsexeeagent.out`
- `xsuaaserver.out`

The additional XS advanced services "HANA Broker" and "Instance Manager" write logs to the following files:

- `xshanabroker_0.log`
- `xsinstancemanager_0.log`

Audit Logs

Each of the platform components writes detailed audit logs, which contain information about login attempts and any changes made to resources managed by the platform. The audit logs are stored in the following files in the SAP HANA trace directory on the respective SAP HANA host:

- `dbxscontroller_audit_0.log`
- `hdbxsexecutionagent_audit_0.log`
- `uaa-audit.log`

i Note

Audit logs are never deleted automatically. For more information about deleting audit logs manually, see *Data Protection and Privacy in Related Information*.

Access Logs

For each route, the platform router maintains access log files with the naming pattern `access_log-<route guid>-<port>.log`, as illustrated in the following example

```
access_log-229dba84-8ed6-45d2-91b6-7b142cc58177-51004.log
```

To display details of a specific route's globally unique identifier (guid) as well as information about any bound applications, you can use the command `xs routes --guids`, as illustrated in the following example:

Output Code

```
$ xs routes --guids

Getting routes in org "orgname" / space "SAP" as XSA_ADMIN...

host domain      port path type apps      guid
-----
-
  example.org 51006 / HTTP auditlog-server
aldac8fa-620a-4093-8...
  example.org 51007 / HTTP auditlog-broker
66eefa83-99fe-4d42-9...
  example.org 51008 / HTTP deploy-service 743db7a4-cdd6-43f0-
a...
  example.org 51009 / HTTP product-installer
a944c3e9-2367-4c0b-8...
[...]
```

→ Tip

Access logs to applications can also be displayed with the command `xs logs`. For more information, see *Related Information*.

To display access logs for the XS Controller and XS UAA, use the following files:

- `access_log-controller-route-30030.log`
- `access_log-external-uaa-route-30032.log`

The following example shows what an individual entry looks like in an application access log:

Output Code

Access Log Entries

```
[13/Feb/2018:13:23:50 +0100] 7.7.007.007 - - to example.org:30030 "GET /v2/
info HTTP/1.1" 200 sent 471 in 182 by 000-controller-instance
```

An access log line has the following format:

- The timestamp at which the request was finished
- The remote IP address
- The client identification if available or "-" if not available
- The user name in case of basic authentication or "-" if not available
- The target host and port
- The http request
- The response code
- The number of bytes transferred
- The request processing time in milliseconds
- The target system identifier

In the example above, a request was sent from the IP address 7.7.007.007 to target "example.org", requesting the resource "GET /v2/info HTTP/1.1". The response had HTTP status code 200 and 421 bytes were sent in 182 milliseconds by the XS Controller.

Related Information

[Data Protection and Privacy Tools in XS Advanced \[page 1752\]](#)

12.2.8 Platform Sizing in XS Advanced

Set up usage profiles and resource consumption at the application and platform levels.

According to the business scenario and load profile, XS advanced platform services have a resource consumption that needs accurate alignment with the available hardware resources. For example, operation of the XS advanced platform requires certain amount of disk space, main memory and number of network connections. In addition, the resource consumption it highly depends on the applications deployed on the platform.

i Note

A tradeoff is required between the platform's resource consumption and the maximum load the server can handle.

By default, XS advanced is optimized for low-memory consumption, which initially limits the maximum number of concurrent requests sent by platform and application users. If you want to configure XS advanced to handle higher loads, you may need to resize the platform services accordingly without hitting resource limits. Usage profiles help administrators to cope with this task as described in the following sections.

Usage Types and Profiles

To simplify sizing configuration, XS advanced comes with a predefined set of usage and sizing profiles enabling administrators to configure the server with a best-fit usage type and size profile (a T-shirt size). Starting with version SPS03, XS advanced server provides the following usage types:

- `PlatformUsage` (Platform Usage)
- `AppUsage` (Application Usage)

Platform Usage

XS advanced systems configured with the usage type `PlatformUsage` are primarily intended to serve users that interact with the platform interface (`xscontroller`), for example, when pushing applications, viewing the application status, or administrating global platform settings. For this reason, the chosen size profile scales the platform services that are responsible for processing such requests. However, the platform router has a default (medium) layout that is independent of the size profile and which assumes that only a few requests to application end points are expected. The typical use case for this usage setting is a development system, and the impact of the profile size on the platform service capabilities is illustrated in following table.

Platform Usage Profile Settings

Profile Size (Short)	Profile Size (Long)	Max. Concurrent Platform Requests
S	Small	4
M	Medium	16
L	Large	64
XL	Extra Large	256

Application Usage (AppUsage)

Systems that are mainly intended to process application requests from business users should be configured with usage type `AppUsage`. In this scenario, the server infrastructure for handling platform user requests has a default layout while the platform router is scaled according to the chosen size profile. A production system would typically use this setting. The impact of the profile size is illustrated in following table.

Application Usage Profile Settings

Profile Size (Short)	Profile Size (Long)	Max. Concurrent Application Requests
S	Small	100
M	Medium	2,000
L	Large	8,000
XL	Extra Large	32,000

Server Configuration

By default, all usage types have profile size M (medium). You may adapt the profile size of a dedicated usage type in `global.ini` by means of the property `xsa_sizing` in section `system_information`. Using the format `<usage>:<size>` you can set the size of a specific usage type to one of the pre-defined sizes (L, XL, etc.). The following example shows a server configuration optimized to serve a maximum of 64 platform requests in parallel:

```
[system_information]
xsa_sizing = PlatformUsage:L
```

The following example shows a server configuration optimized to handle a maximum of 8,000 concurrent application requests:

```
[system_information]
xsa_sizing = AppUsage:L
```

It is possible to specify size profiles for all different usage types individually, which could be helpful in mixed scenarios. In the example below the server is configured to handle a very large number of platform users and a

large number of business requests at the same time, which is the kind of scenario you would expect in a development and test system running extensive application tests:

```
[system_information]
xsa_sizing = PlatformUsage:XL, AppUsage:L
```

Providing the usage type is optional, and if it is not specified, the profile size applies to all usage types by default. So 'xsa_sizing = L' induces the same server configuration as 'xsa_sizing = PlatformUsage:L, AppUsage:L'. This ensures the compatibility of sizing configurations for pre-SPS03 released versions of SAP HANA.

i Note

Changes to the `xsa_sizing` property require a restart of the XS advanced server (for example, with `xsa restart`) in order to make SAP HANA aware of the modified configuration.

Resource Consumption

Depending on the chosen platform sizing configuration, the platform services have a certain resource consumption that needs to be taken into consideration when XS advanced is activated within an SAP HANA system. For example, the XS advanced platform has a significant impact on the following resources:

- Disk space and file IO
- Main memory
- Network connections and network IO

The profile size mainly influences the amount of occupied main memory, the maximum number of network connections, and the maximum file and network IO rates. Disk space, however, is not correlated to the profile size; it is associated with the number and size of the deployed applications.

→ Tip

For more information about recommendations for resource-consumption settings, see *Related Information*.

Application Sizing

The usage types and profile sizes specified in the corresponding section of the configuration-parameters (`.ini`) file enable only the scaling of the core platform services, which do not affect applications that are deployed to the XS advanced platform (including system applications such as the deploy service or audit log service). To meet load requirements, applications generally need to be scaled horizontally or vertically. For more information, see *Related Information*.

Related Information

[SAP Note 2618752 \(Resource Consumption in XS Advanced\)](#) 

[Scaling Applications in XS Advanced \[page 1690\]](#)

[XS Advanced System Configuration Parameters \[page 1824\]](#)

12.2.9 Configuring the XS Advanced Platform Router

Use a template or INI parameters to configure the XS advanced platform router.

XS advanced model uses a dedicated instance of the SAP Web Dispatcher as the central platform router. The configuration of this Web Dispatcher instance is managed by the XS Controller and is highly dynamic; the configuration regularly changes, for example, when applications are started, stopped, or scaled, or when routes are mapped to (and unmapped from) applications. Due to this dynamic behavior, it is not possible for the administrator to change the configuration of the Web Dispatcher directly. Instead, changes to the configuration are defined in a configuration template which is used by the XS advanced platform to build the final Web Dispatcher configuration that is used by the XS advanced platform router.

The information in this section describes some examples that demonstrate the process including details of the following steps:

- [Configuring the Platform Router with the Configuration Template \[page 1842\]](#)
- [Configuring the Platform Router with INI Parameters \[page 1843\]](#)

Configuring the Platform Router with the Configuration Template

XS advanced provides a template file `sapwebdisp.template` which you can use to override or add SAP Web Dispatcher configuration properties. For a complete reference of all configuration parameters, see the SAP Web Dispatcher documentation in *Related Information* below.

i Note

Properties defined in the template file take precedence over properties determined by the XS advanced platform.

Saving any changes made to the configuration template file prompts XS advanced to automatically update the Web Dispatcher, which means that it is not necessary to restart the Web Dispatcher manually. The template file can be edited by the `<sid>adm` user and is located in the "xs/" directory of the SAP HANA installation, as illustrated in the following path:

```
/hana/shared/<SID>/xs/controller_data/controller/router/webdispatcher/conf/  
sapwebdisp.template
```

Any changes made to the XS advanced router-configuration template `sapwebdisp.template` will be reflected in the actual Web Dispatcher configuration file named `sapwebdisp.pfl`

i Note

Some configuration properties of the SAP Web Dispatcher cannot be changed dynamically.

For changes to configuration properties of the SAP Web Dispatcher that cannot be updated dynamically, refer to the SAP Web Dispatcher documentation, for example, in *Related Information* below. In these cases, a restart

of the SAP Web Dispatcher is required, for example, by sending signal 3 (SIGQUIT) to “kill” the Web Dispatcher process directly. Alternatively, you can restart the XS Controller service or the entire XS advanced platform, for example, using the `xsa restart` command. Bear in mind that restarting the XS Controller or the XS advanced platform will cause a longer downtime than restarting only the Web Dispatcher.

⚠ Caution

It is not recommended to use the template file to override any parameters that include an index in their key.

Although it is possible to override parameters that include an index in their key (for example, `wdisp/system_xx`), it is not recommended. This is because the XS advanced run time cannot guarantee that these indexes remain stable and always reference the same logical system. XS advanced incrementally increases the index for the parameters it sets, starting at zero. If you want to add additional parameters that use an index in their key, make sure you use a very high value as the index.

Configuring a Load Balancing Algorithm

The SAP Web Dispatcher provides several options for requesting load balancing. For more details about what those options are and how to set up and use them, see *Load balancing via SAP Web Dispatcher* in *Related Information* below.

Serving XS Classic with the XS Advanced Web Dispatcher

One use case for adding additional configuration to the XS advanced Web Dispatcher is related to XS classic, for example, to enable you to serve XS classic and XS advanced applications on the same TCP port. This is useful when running XS advanced in hostname-based routing mode, as you only need to open a single port for HTTPs communication in that case. The SSL settings will be taken from the XS advanced Web Dispatcher, which means you only need to configure the HTTP endpoint of XS classic, so that internal communication between the XS advanced Web Dispatcher and the XS classic Web Dispatcher can be established. To enable such a configuration, add the following line to your `sapwebdisp.template` file:

≡ Sample Code

Connect the XS Advanced and XS Classic Web Dispatchers in `sapwebdisp.template`

```
wdisp/system_999=NAME=ZZZ, SID=ZZZ, SRCURL=/, SRCVHOST=xsc.<default-domain>:<router-port>, SSL_ENCRYPT=0, EXTSRV=<xsc-classic-http-url>
```

Values for the placeholders `<default-domain>` and `<router-port>` can be obtained from your XS advanced, system-parameter configuration, which is described in *XS Advanced System Configuration Parameters* in *Related Information*. The value of `<xsc-classic-http-url>` `http://127.0.0.1:80<instance-nr>`". After you have adapted the template file, the XS classic start page should be available at " should reflect the HTTP end point of the XS classic Web Dispatcher. If both XS advanced and XS classic are running on the same host (for example, within a single-host SAP HANA installation), the end point for the Web Dispatcher could be `https://xsc.<default-domain>:<router-port>`".

Configuring the Platform Router with INI Parameters

Some settings of the XS advanced Web Dispatcher can be configured by means of parameters within the file `xscontroller.ini` as described in *XS Advanced System Configuration Parameters* in *Related Information*. For

example, the property `router_https` determines whether the Web Dispatcher exposes HTTP or HTTPS endpoints.

Configuring TLS Versions and Cipher Suites

The TLS settings of the Web Dispatcher server ports used by the XS advanced platform can be configured by means of an INI parameter. The following INI settings are available in the `xscontroller.ini` file. If you are running a multi-host system, with multiple `xs_worker` roles configured, the settings also need to be added to the `xsexecagent.ini` file. This is because an additional Web Dispatcher is started by Execution Agents, which are not running on the XS advanced master host.

INI Parameter for Router Cipher-Suite Setting

Property Name	Section	Default	Description
<code>Router.WebDispatcher.CipherSuites</code>	<code>router</code>	<code>135:PFS:HIGH::EC_P256:EC_HIGH</code>	should reflect the HTTP endThe TLS cipher suite settings, provided on server ports exposed by the Web Dispatcher (enables TLS 1.0, TLS 1.1 and TLS 1.2)

For example to disable all TLS versions below TLS 1.2, but retain all other cipher settings, change the parameter in the INI file `xscontroller.ini` as illustrated in the following example:

```
Router.WebDispatcher.CipherSuites = 519:PFS:HIGH::EC_P256:EC_HIGH
```

For more details about the format required for the cipher-suite parameters, see SAP Note [510007](#).

Note

All internal HTTPS ports used by XS advanced are restricted to TLS 1.2 by default.

Related Information

[XS Advanced System Configuration Parameters \[page 1824\]](#)

[SAP Web Dispatcher](#)

[Load balancing via SAP Web Dispatcher \(SAP Community Wiki\)](#)

12.2.10 Maintaining Single Sign-On for XS Advanced Applications

You can configure XS advanced applications to use single sign-on (SSO) authentication to confirm the logon credentials of a user calling an application service.

Single sign-on is a convenient way to authenticate a user against the XSUAA service using one or multiple available certificates. The XS advanced runtime enables you to use X.509 authentication as well as SPNEGO and Kerberos.

i Note

To prevent a login-logout loop with SSO, it is recommended to configure a custom logout page for the XS advanced application. For more information, see the reference documentation for the application router configuration syntax.

Related Information

[Configure SSO with X.509 Authentication for XS Advanced Applications \[page 1845\]](#)

12.2.10.1 Configure SSO with X.509 Authentication for XS Advanced Applications

Enable SSO with X.509 certificates for user logon to applications in the XS advanced run-time environment.

Prerequisites

- You have the role `USER_ADMIN` on your SAP HANA database.
- You have the role `XS_CONTROLLER_ADMIN` on your XS advanced runtime.
- You have an SAP HANA administration tool.
- You have the root certificate of the X.509 user certificates.
- The parameters `uaa.oidc.enableoidc` and `uaa.oidc.enablex509` in the `xsuaaserver.ini` file are set to `true`.

Procedure

1. Add the trust certificate for client authentication.

Add the root certificate of the X.509 user certificates in PEM format to your XS advanced runtime platform.

Use the `xs trust-certificate` command with the `--client-auth` option, as illustrated in the following example:

```
xs trust-certificate <ALIAS> --client-auth -c <PATH>
```

You see a confirmation in your command prompt.

```
Adding trusted certificate <ALIAS> as <user>...  
OK
```

2. Create a database user whose identity is defined in an X.509 certificate issued by your certificate authority (CA).
 - a. Create a new user in the SAP HANA database based on the details specified in an existing X.509 certificate.

The following example shows how to use the SQL statement `CREATE USER WITH IDENTITY` to create the database user “MyUserName” and the corresponding X.509 certificate mapping:

```
CREATE USER MyUserName WITH IDENTITY 'CN=MyUserName, O=SAP-AG, C=DE'  
ISSUER 'CN=SSO_CA, O=SAP-AG, C=DE' FOR X509
```

- b. Configure the user in SAP HANA.

To edit the user for whom you want to enable SSO, use the SAP HANA administration tool of your choice: select [X509](#) and [Configure X.509 user certificates](#).
3. Use a Web browser to test the logon authentication settings for the XS advanced application.

When you enter the URL for your application in the Web browser, the Web browser prompts you to select a certificate, which enables you to log on without supplying logon credentials manually.

Related Information

[Maintaining Trust Certificates in XS Advanced \[page 1706\]](#)

12.2.10.2 Configure SSO with SPNEGO and Kerberos for XS Advanced Applications

Enable SSO with SPNEGO and Kerberos for user logon to applications in the XS advanced run-time environment.

Prerequisites

- You have configured Kerberos and SPNEGO support in your SAP HANA database.
- You have an SAP HANA administration tool.

Context

In addition to configuring Kerberos in SAP HANA, see SAP Notes [1813724](#) and [1837331](#).

i Note

For SPNEGO, the Kerberos keytab must contain an entry for Service Principal Name (SPN): `HTTP/<fully-qualified-host-name>`

Procedure

1. Open your SAP HANA administration tool and set the parameter `uaa.oidc.enablespnego` in the `xsuaaserver.ini` file to `true`.
2. Configure your SAP HANA user.

Edit the user in your SAP HANA administration tool: select *Kerberos* and enter an *External ID*.

3. Test your configuration.

With the Python script from SAP Note [1813724](#) you may verify, that your logon to SAP HANA with your Kerberos credentials works.

Related Information

[Configure Kerberos for SAP HANA Database Hosts \[page 730\]](#)

13 Data Access

SAP HANA supports the integration of data from many data sources to enrich your applications and deliver in-depth analysis. These include federated queries, data replication, remote data sync, and processes to improve data quality.

This section provides an overview of the tools and technologies that are available with SAP HANA or supported by SAP HANA for data access and data virtualization. For more information about the administration of these technologies, as well as other operations topics, refer to the documentation indicated.

Native Capabilities for Data Access, Integration, and Quality in SAP HANA

Capability	Description	More Information
Data Federation with SAP HANA Smart Data Access (SDA)	SAP HANA SDA enables you to create virtual tables in SAP HANA that point to virtual tables on remote sources, such as SAP ASE, SAP IQ, Hadoop, and Teradata.	See SAP HANA Smart Data Access [page 1849] and SAP HANA Hadoop Integration [page 1962]
Data Replication and Transformation	SAP HANA smart data integration provides the architecture that supports all types of data delivery in SAP HANA: real-time, batch, and federation (SDA). It includes both data replication and data transformation services.	See the documentation for SAP HANA smart data integration and SAP HANA smart data quality option on SAP Help Portal at https://help.sap.com/viewer/p/HANA_SMART_DATA_INTEGRATION .
Remote Data Synchronization	SAP HANA Remote Data Sync is a session-based synchronization technology designed to synchronize SAP SQL Anywhere remote databases with a consolidated database.	See the documentation for SAP HANA real-time replication on SAP Help Portal at http://help.sap.com/viewer/p/SAP_HANA_REAL_TIME_REPLICATION .

Data Replication Technologies in the Extended SAP HANA Landscape

Capability	Description	More Information
Trigger-Based Replication	The trigger-based replication method uses the SAP Landscape Transformation (LT) Replication Server component to pass data from the source system to the SAP HANA database target system.	See the documentation for SAP HANA real-time replication on SAP Help Portal at http://help.sap.com/viewer/p/SAP_HANA_REAL_TIME_REPLICATION .

Capability	Description	More Information
SAP HANA Direct Extractor Connection	SAP HANA Direct Extractor Connection (DXC) provides SAP HANA with out-of-the-box foundational data models based on SAP Business Suite entities, and is a data acquisition method as well.	See SAP HANA Direct Extractor Connection Implementation Guide
Extraction Transformation Load-Based Replication	Extraction Transformation Load (ETL)-based data replication uses SAP Data Services (also called Data Services) to load relevant business data from SAP ERP to the SAP HANA database. This lets you read the business data on the application layer level.	See the documentation for SAP Data Services on SAP Help Portal at http://help.sap.com/viewer/p/SAP_DATA_SERVICES
Log-Based Replication	SAP Replication Server (SRS) moves and synchronizes transactional data including DML and DDL across the enterprise, providing low impact, guaranteed data delivery, real-time business intelligence, and zero operational downtime.	See the documentation for SAP Replication Server on SAP Help Portal at http://help.sap.com/viewer/p/SAP_REPLICATION_SERVER

13.1 SAP HANA Smart Data Access

SAP HANA smart data access allows you to access remote data as if the data was stored in local tables in SAP HANA, without copying the data into SAP HANA.

This capability provides operational and cost benefits and supports the development and deployment of next-generation analytical applications requiring the ability to access, synthesize, and integrate data from multiple systems in real-time regardless of where the data is located or what systems are generating it.

In SAP HANA, you use linked database or create virtual tables, which point to remote tables in different data sources and then write SQL queries in SAP HANA, using these virtual tables. The SAP HANA query processor optimizes these queries by executing the relevant part of the query in the target database, returning the results of the query to SAP HANA, and then completing the operation. Physical data movement is not supported by SAP HANA smart data access.

As part of the HANA core system, no additional licensing is required to use smart data access. However, additional support packages and patches are available for download from the SAP Software Download Center, and are installed using the SAP HANA database lifecycle manager (HDBLCM).

SAP HANA and the remote source must be using the same hardware platform. For example, if your Oracle database is running on Intel-based hardware, and your SAP HANA system is using an IBM Power system, you can't create an Oracle remote source. See [2600176 - Smart Data Access - Supported Remote Source Databases and Versions](#) for a list of supported remote source databases and versions.

13.1.1 Setting Up ODBC Drivers

The communication between SAP HANA and a remote data source is based on the ODBC protocol. To use the protocol, install the appropriate drivers for the databases you want to connect to using SAP HANA smart data access.

Supported Remote Source Databases

See [2600176 - Smart Data Access - Supported Remote Source Databases and Versions](#) for a list of supported remote source databases and versions.

.ODBC.INI File

SAP HANA smart data access requires that an `.odbc.ini` file exist in the administrator's home directory, even if empty. If using DSN references to create the remote source, create one entry in the `.odbc.ini` file for each remote source.

ODBC Driver Installation Location

Install ODBC driver library files in a location that is searched by the SAP HANA server. Libraries installed in the SAP HANA `exe` directory are found automatically. Libraries installed elsewhere require an entry in the `LD_LIBRARY_PATH` environment variable to point to this location. If SAP HANA is unable to locate the libraries, you may experience messages during SAP HANA smart data access queries stating the driver could not be loaded. To prevent possible interference with software maintenance of the SAP HANA product, we recommend that you install the ODBC files to a location other than the SAP HANA `exe` directory and add the path to the `LD_LIBRARY_PATH` environment variable.

In a scale-out landscape, install the driver on all hosts.

The `LD_LIBRARY_PATH` environment variable is configured by creating or modifying the `.customer.sh` file in the home directory of the SAP HANA administrator user to include the following entry:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/<path_to_driver_directory>
```

Since the `LD_LIBRARY_PATH` environment variable includes the directory path, the `'Driver='` entry in the DSN section of the `odbc.ini`, or in the `CREATE REMOTE SOURCE` command only needs to contain the library name. Restart the SAP HANA system to apply the change to the `.customer.sh` file. To validate that the changes in `.customer.sh` have taken effect, logged in as the `<sid>adm`, execute:

```
echo $LD_LIBRARY_PATH
```

❖ Example

If the IQ ODBC libraries are installed in `/usr/sap/sapiq/IQ-16_0/lib64`, the administrator's `$HOME/.customer.sh` file should include:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/sap/sapiq/IQ-16_0/lib64
```

This statement adds the directory `/opt/sap/IQ-16_0/lib64` to the end of the search path SAP HANA uses to find libraries.

HADOOP (ODBC) Driver

The HADOOP (ODBC) driver is a Hive ODBC adapter, which requires installation of the Apache Hive ODBC driver on the SAP HANA server. For more information, see [SAP HANA Hadoop Integration > Using the Simba ODBC Driver to Connect to Hive](#).

ODBC Driver Software Ownership

We recommend that you do not make the `<sid>adm` user the software the owner of the ODBC driver library files. Create a new user to become the software owner.

Remote Data Sources

The process to install, configure, and create remote data sources depends upon the ODBC driver. Refer to the specific driver for details.

13.1.1.1 SAP HANA ODBC Driver

SAP HANA Smart Data Access uses the SAP HANA ODBC driver installed with the SAP HANA server.

The SAP HANA ODBC driver only supports access to a remote SAP HANA database version SPS 10 or later. No additional configuration of the driver is required.

13.1.1.2 SAP IQ ODBC Driver

The SAP IQ ODBC driver is included in the SAP HANA smart data access maintenance software component, which is available for download from the SAP Software Download Center.

Context

⚠ Caution

Do not use the SAP IQ ODBC driver included in the IQ Network Client for Linux package. Unexpected behavior has been reported when using this driver with smart data access.

Procedure

1. Download and extract the maintenance SAR file to a temporary location on the SAP HANA host.
2. Log on to the SAP HANA host as the SAP HANA software owner (<sid>adm), and change to the `hdblcm` directory within the <SID> folder.
3. Start the installer. Be sure to include the clause to the location of extracted smart data access package:

```
./hdblcm --component_dirs=/path to package
```

4. When installation is complete, use the `find` command to verify that the ODBC driver is installed. By default, the ODBC driver (`libdbodbc17_r.so` library file) is installed to `/hana/shared/<SID>/federation`. Note the full path as it is needed in future steps.

```
find /hana/shared -name libdbodbc17_r.so
```

5. Switch to the <sid>adm \$HOME directory.
6. Create or modify the `.customer.sh` file to add the following line:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:<odbc_path>
```

where <odbc_path> is the path to the ODBC driver.

7. Restart the SAP HANA system to apply the change to the `.customer.sh` file.

Related Information

[SAP Software Download Center](#) 

13.1.1.3 SAP Adaptive Server Enterprise ODBC Driver

The SAP Adaptive Server Enterprise (ASE) ODBC driver is included in of the SAP HANA smart data access maintenance software component, which is available for download from the SAP Software Download Center.

Procedure

1. Download and extract the maintenance SAR file to a temporary location on the SAP HANA host.
2. Log on to the SAP HANA host as the SAP HANA software owner (<sid>adm), and change to the `hdblcm` directory within the <SID> folder.
3. Start the installer. Be sure to include the clause to the location of extracted smart data access package:

```
./hdblcm --component_dirs=/path to package
```

4. When installation is complete, use the `find` command to verify that the ODBC driver is installed. By default, the ODBC driver (`libsydbrvodb-sql11en8.so`) is installed to `/hana/shared/DT1/federation`. Note the full path as it is needed in future steps.

```
find /hana/shared -name libsydbrvodb-sql11en8.so
```

5. Switch to the <sid>adm \$HOME directory.
6. Create or modify the `.customer.sh` file to add the following line:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:<odbc_path>
```

where <odbc_path> is the path to the ODBC driver.

7. Restart the SAP HANA system to apply the change to the `.customer.sh` file.

Related Information

[SAP Software Download Center](#) 

13.1.1.4 SAP Event Stream Processor ODBC Driver

SAP Event Stream Processor (ESP) on Linux ODBC driver is included in the ESP software package, and is available for download from the SAP Software Download Center.

Prerequisites

Java is installed on the SAP HANA host.

Context

The ESP ODBC driver can be installed using the *Custom* option in the full ESP installer, or from the `.../archives/esp_odbc` folder where you extracted the ESP package. If Java is not installed, and you use the full installer, the installation appears to run successfully, but the ODBC driver is not installed. If you install from the `esp_odbc` folder, an error message appears and the installation fails if Java is not installed.

Procedure

1. Download and extract the ESP package to a temporary location on the SAP HANA host.
2. Download and install the unixODBC driver manager, which is included in the SAP Event Stream Processor installation package or available from your Linux provider.
3. Logged on to the SAP HANA host as the SAP HANA software owner (`<sid>adm`), change to the location where you extracted the ESP package.
4. Change to the `.../archives/esp_odbc` folder and type:

```
./setup.bin
```

5. When installation is complete, use the `find` command to verify that the ODBC driver is installed. By default, the ODBC driver (`streamingpsqlodbc_lib.so`) is installed to `/opt/sybase/lib`. Note the full path as it is needed in future steps.

```
find /opt/sybase/lib -name libstreamingpsqlodbc_lib.so
```

6. Switch to the `<sid>adm` \$HOME directory.
7. Create or modify the `.customer.sh` file to add the following lines:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:<odbc_path>/lib
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:<odbc_path>/STREAMING-1_0/odbc
```

where `<odbc_path>` is the path to the ODBC driver.

8. Restart the SAP HANA system to apply the change to the `.customer.sh` file.

Related Information

[SAP Software Download Center](#) 

13.1.1.5 SAP HANA Streaming Analytics ODBC Driver

SAP HANA streaming analytics ODBC driver is installed as part of the streaming analytics installation or the streaming analytics client, and is available for download from the SAP Software Download Center.

Context

You can create a remote source to a streaming analytics system running on another host or on the current host. If the remote source points to the current host, the streaming analytics ODBC driver is already installed as part of the streaming analytics instance. If the remote source points to a different host, install the streaming analytics client. The version of the client must be the same as the current host, not the remote host. For example, the current host is running SAP HANA 2.0 SPS 00, while the remote host is running SAP HANA 1.0 SPS 12. You would install the 2.0 client, not the SPS 12 client.

Procedure

1. If required, install the SAP HANA streaming analytics client, which includes the ODBC driver.
2. Download and install the unixODBC driver manager, which is included in the streaming analytics installation package or available from your Linux provider.
3. Log on to the SAP HANA host as the SAP HANA software owner (<sid>adm).
4. Switch to the <sid>adm \$HOME directory.
5. Create or modify the `.customer.sh` file to add the following lines.

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$STREAMING_HOME/odbc
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$STREAMING_HOME/lib
```

i Note

If streaming analytics is already installed, these entries may already exist.

6. Restart the SAP HANA instance to apply the changes to the `.customer.sh` file.

Related Information

[SAP Software Download Center](#) 

13.1.1.6 SAP MaxDB ODBC Driver

The SAP MaxDB ODBC driver is included in the SAP MaxDB LinuxDB server installation package, which is available for download from the SAP Software Download Center.

Procedure

1. Download and extract the client installation files to a temporary location on the SAP HANA host.
2. On the SAP HANA host, logged on as the SAP HANA software owner (<sid>adm), switch to the location of the extracted files.
3. Type:

```
./SDBINST
```

4. When prompted for the software components to install, select *ODBC only*.
5. When installation is complete, use the `find` command to verify that the ODBC driver is installed. By default the ODBC driver (`libsdbodbcw.so` library file) is installed to `/sapdb/clients/MAXDB`.

```
find /sapdb/ -name libsdbodbcw.so
```

6. Change to the <sid>adm \$HOME directory.
7. Add the following line to the `.customer.sh` file.

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:<odbc_path>
```

where <odbc_path> is the path to the ODBC driver.

8. Restart the SAP HANA system to apply the change to the `.customer.sh` file.

Related Information

[SAP Software Download Center](#) 

13.1.1.7 Teradata ODBC Driver

Teradata drivers are packaged and distributed as a Linux RPM file. You can obtain these RPMs from Teradata.

Procedure

1. Download and extract the files to a temporary location on the SAP HANA host.

2. On the SAP HANA host, logged on as the SAP HANA software owner (<sid>adm), install the extracted files.
 - For version 13, installation of TeraGSS_redhatlinux-i386__linux_i386.13.10.00.06-1.tar.gz normally fails. While not required for the ODBC installation, the installation attempt prevents `sudo rpm -i tdodbc-13.10.00.04-1.noarch.rpm` from failing with missing dependencies, which is required as part of the driver installation.
 - For version 14, run `sudo rpm -i tdicu-14.10.00.04-1.noarch` to prevent a dependency error.
3. Change the default Kerberos 5 setup.

The Teradata driver loads GSS API libraries from the OS folders, which conflict with the version of libraries loaded by SAP HANA during installation. Since SAP HANA does not support single sign-on for Teradata remote sources, you can safely disable the Kerberos 5 mechanism.

- a. Edit the `/<installaton_path>/teragss/site/TdgssUserConfigFile.xml` file and add:

```
<Mechanism Name="KRB5">
  <MechanismProperties MechanismEnabled="no" />
</Mechanism>
```

- b. Remove the `/<installaton_path>/teragss/site/linux-x8664/<version>/TdgssUserConfigFile.xml` file, if it exists.
- c. As `sudo rmp`, execute:

```
/opt/teradata/teragss/linux-x8664/client/bin/run_tdgssconfig
```

13.1.1.8 SQL Server ODBC Driver

SAP HANA remote data sources require the Microsoft ODBC driver for SQL Server.

Procedure

1. Download and extract the Microsoft ODBC driver package to a temporary location on the SAP HANA host.
2. Download and install the unixODBC driver manger as specified by SQL Server.
3. On the SAP HANA host, logged on as the SAP HANA software owner (<sid>adm), install the Microsoft ODBC driver.
The driver installation creates the following directory: `/opt/microsoft/msodbcsql/lib64`.
4. Change to the `<sid>adm $HOME` directory.
5. Add the following line to the `.customer.sh` file:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:<odbc_path>
```

where `<odbc_path>` is the path to the ODBC driver.

Depending on your setup, you might also need to add `/usr/local/lib64` to the SAP HANA administrator user's `PATH` variable.

6. Restart the SAP HANA system to apply the change to the `.customer.sh` file.

13.1.1.9 IBM DB2 Driver

IBM DB2 drivers are available on the IBM website.

Procedure

1. Download and extract the IBM DB2 ODBC package to a temporary location on the SAP HANA host.
2. Download and install the unixODBC driver manager as specified by IBM DB2.
3. Log on to the SAP HANA system as the root user.
4. Create the following directory: `/opt/ibm/db2` and extract the ODBC package there.
5. Change to directory: `/opt/ibm/db2/odbc_cli/clidriver/cfg`.
6. Add a new entry to the `db2cli.ini` file, replacing the applicable information for your system.

Sample Code

```
[TEST_DB2]
Database=<database_name>
Protocol=TCPIP
Port=50010
Hostname=<machine_name>
```

7. Use the `find` command to verify the ODBC driver is installed. By default, the ODBC driver (`libdb2o.so.1` library file) is installed to `/opt/ibm/db2/odbc_cli/clidriver/lib`. Note the full path as it is needed in future steps.

```
find /opt/ibm -name libdb2o.so.1
```

8. Log on as the `<sid>adm` user and switch to the `<sid>adm $HOME` directory.
9. Create or modify the `.customer.sh` file to add the following line:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:<odbc_path>
```

where `<odbc_path>` is the path to the ODBC driver.

10. Restart the SAP HANA system to apply the change to the `.customer.sh` file.

13.1.1.10 Oracle Database ODBC Driver

SAP HANA remote data sources for Oracle databases require the ODBC driver for Oracle.

Procedure

1. Download and extract the Oracle ODBC driver package to a temporary location on the SAP HANA host.

2. Download and install the unixODBC driver manager as specified by Oracle.
3. Install the Oracle ODBC driver package.
4. Change to `<install_path>/client64/lib` and verify it contains the ODBC driver `libsqora.so.<version>`.
5. Switch to the SAP HANA software owner (`<sid>adm`), then change to the `<sid>adm $HOME` directory.
6. Create a new `tnsnames.ora` file in the home folder. For details on the `tnsnames.ora` file, refer to Oracle documentation.
7. Create or modify the `.customer.sh` file to add the following lines:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:<install_path>/client64/lib
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib
export TNS_ADMIN=~/
```

8. Restart the SAP HANA system to apply the change to the `.customer.sh` file.

13.1.1.11 IBM Netezza Driver

IBM Netezza drivers are available on the IBM website.

Procedure

1. Download and extract the IBM Netezza ODBC package to a temporary location on the SAP HANA host.
2. Download and install the unixODBC driver manager as specified by IBM Netezza.
3. Log on to the SAP HANA system as the root user.
4. Create a directory: `/usr/local/nz` and extract the ODBC package.
5. Change to the `linux64` folder within the path of the extracted package and type:

```
./unpack
```

6. Use the `find` command to verify the ODBC driver is installed. By default, the ODBC driver (`libnzodbc.so` library file) is installed to `usr/local/nz/lib64`. Note the full path as it is needed in future steps.

```
find /usr/local/nz -name libnzodbc.so
```

7. Log on as the `<sid>adm` user and switch to the `$HOME` directory.
8. Create or modify the `.customer.sh` file to add the following line:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:<odbc_path>
```

where `<odbc_path>` is the path to the ODBC driver.

9. Restart the SAP HANA system to apply the change to the `.customer.sh` file.

13.1.1.12 Google BigQuery ODBC Driver

SAP HANA remote data sources for Google BigQuery databases require the ODBC driver for Google BigQuery database.

Procedure

1. Download the Linux 64-bit Simba ODBC driver for Google BigQuery from the Google Cloud Platform site.
2. Download and install the unixODBC driver manager as specified by Google BigQuery.
3. On the SAP HANA host, logged on as the SAP HANA software owner (<sid>adm), extract the file using the `--directory=`/`<install_path>` option.
4. Change to the folder you extracted the file to and edit the `simba.googlebigqueryodbc.ini` file. Change the following parameters:
 - Set `DriverManagerEncoding` to use UTF-16. You will be unable to connect to the Google BigQuery database if this parameter is not set.
 - Set `ErrorMessagePath` to point to `/`<install_path>/`simba/googlebigqueryodbc/ErrorMessage`s
 - In the Generic `ODBCInstLib` section, comment out `ODBCInstLib=libiodbcinst.so`
 - In the `unixODBC` section uncomment `ODBCInstLib=libodbcinst.so.2`.
5. Create or modify the `customer.sh` file to add the following lines:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/$install_path/simba/googlebigqueryodbc/lib/64
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:<path_to_unixODBC_library>
```

6. Restart the SAP HANA system to apply the change to the `.customer.sh` file.

13.1.2 Remote Source Credential Types

Remote sources support two types of credential modes to access a remote source: technical user and secondary credentials.

Credential Type	Description
Technical User	A valid user and password in the remote database. This valid user is used by anyone using the remote source.
Secondary Credentials (Preferred option)	A unique access credential on the remote source assigned to a specific user. See <i>Managing Secondary Credentials</i> .

Both credential types can be set up with either a password or Kerberos single sign-on.

Related Information

[Managing Single Sign-On \(SSO\) with Kerberos \[page 1915\]](#)

13.1.3 Linked Database Overview

Linked database allows DML queries on remote data sources without the need to first create virtual tables for each table referenced in a query before executing the query.

This makes ad-hoc access to remote data much more convenient.

Linked database uses a three-part namespace to directly identify the remote database, schema, and table name. The three-part name syntax is defined as:

```
<remote_source>.<remote_schema>.<remote_table>
```

Queries between tenant databases, known as cross-database access, also use this three-part namespace. It is possible for the name of a remote source and a tenant database to be the same. If this occurs, the database name resolves to the tenant database as a cross-database query, not the remote source.

The linked database feature supports all smart data access remote data sources and requires the LINKED DATABASE system privilege to use the feature.

When a query is executed, the system locates the referenced remote table and then checks to see if it has been used in previous queries. If it hasn't, the system automatically creates an internal virtual table, which is used to process the query. Subsequent queries to the same remote table are automatically redirected to the internal virtual table. Users cannot directly access these internal virtual tables. However, they can be administered for refreshing metadata and housekeeping.

Linked Database Optimized Mode

In SAP HANA SPS 03, linked database introduces the `optimized mode` feature.

i Note

While linked database is supported for all smart data access remote sources, the `optimized mode` feature is only available for SAP HANA to SAP HANA workflows.

With optimized mode enabled, when a query is executed, metadata of referenced remote objects is cached locally. Internal virtual tables are no longer created. This metadata is automatically refreshed. No housekeeping tasks are required. If the remote object referenced is a view, the metadata of the tables, which are part of the view definition, are also cached locally. You don't need to create or refresh statistics for remote objects. They are fetched automatically when needed. The global optimizer may generate a better query plan, resulting in better query performance.

Optimized mode is disabled by default. To enable, when creating a remote source, specify the `linkeddatabase_mode` property. For example:

Sample Code

```
CREATE REMOTE SOURCE MY_HANA_ADAPTER hanaodbc
  CONFIGURATION 'DSN=HANA1;linkeddatabase_mode=optimized'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

For more information on the linked database optimized mode, see [2605574 - Smart Data Access Linked Database Optimized Mode Functional Restrictions](#)

13.1.4 Privilege Maintenance

Various privileges are required to manage remote sources, virtual tables, and linked database.

Remote Source

- Creating a remote source requires the CREATE REMOTE SOURCE system privilege.
- Managing a remote source created by another user requires additional object level privileges on the remote source. No privileges are required to manage your own remote sources.
- Managing other user's credentials on the remote source requires the CREDENTIAL ADMIN system privilege. No privileges are required to manage your own credentials.

Virtual Tables

- Creating virtual tables requires the CREATE VIRTUAL TABLE object level privilege on the remote source.
- Managing virtual tables owned by other users requires additional object level privileges (for example, INSERT, UPDATE, DELETE). No privileges are required to manage your own virtual tables.

Linked Database

Using and managing linked database requires the LINKED DATABASE object level privilege on the remote source, regardless of who created the remote source. This allows ad-hoc linked database access using a three part name (<remote_source>.<schema_name>.<remote_table_name>).

This privilege, validated on the SAP HANA side, controls access to the linked database feature. Granular access control checks are performed on the remote database system to avoid duplication of access control mechanisms. While a technical user can be used for linked database, it is recommended that secondary credentials be used for any linked database operations. Use CREDENTIAL management commands to manage the remote database user performing DML operations.

13.1.5 Creating a Remote Source

Use SAP HANA studio or SAP HANA database explorer to create a remote source. The parameters required vary by adapter.

A remote source can be used to create virtual tables or to use the linked database feature.

For a list of supported databases and versions, see [2600176 - Smart Data Access - Supported Databases and Versions](#).

13.1.5.1 Create an SAP HANA Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to another SAP HANA database.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

An SAP HANA remote source can be used with virtual tables or the linked database feature. An SAP HANA remote source supports failover.

The following syntax examples assume the remote source is configured to use technical user credentials. See *Managing Secondary Credentials* and *Managing Single Sign-On (SSO) with Kerberos* for syntax to use other credential types.

Create DSN Entry

Context

i Note

These steps are only required if you plan to use a DSN entry to create your remote source.

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (<sid>adm), change to the <sid>adm \$HOME directory.
2. Create an .odbc.ini file if it doesn't already exist.

The .odbc.ini must exist, even if empty, regardless of whether a DSN entry is being used.

3. Define one entry in the .odbc.ini file for each remote source. For example:

```
[HANA1]
Driver=libodbcHDB.so
ServerNode=<remote_server_name>:3<remote_instance_number>15
```

4. To enable failover, add the failover server name to the ServerNode property, separated by a comma. For example:

```
[HANA1]
Driver=libodbcHDB.so
ServerNode=<remote_server_name>:
3<remote_instance_number>15,<failover_server_name>:
3<failover_instance_number>15
```

5. To set session specific connection information, add the sessionVariable property. For example:

```
[HANA1]
Driver=libodbcHDB.so
sessionVariable:<session_variable_name>=?
```

6. To enable optimized mode, add the sessionVariable property. For example:

```
[HANA1]
Driver=libodbcHDB.so
linkeddatabase_mode=optimized
```

7. Restart the SAP HANA system to apply the changes to the .INI file.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- If you're planning to create the remote source using a DSN entry, it already exists in the ODBC.ini file.

Procedure

1. In a SQL console, connect to the tenant database.
2. Do one of:
 - If using a DSN entry, execute a CREATE REMOTE SOURCE command, referencing the DSN entry in the .odbc.ini file.

Sample Code

```
CREATE REMOTE SOURCE MY_HANA ADAPTER hanaodbc
  CONFIGURATION 'DSN=HANA1'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

- If not using a DSN entry, execute a CREATE REMOTE SOURCE command specifying all driver properties.

Sample Code

```
CREATE REMOTE SOURCE <HANA1> ADAPTER hanaodbc
  CONFIGURATION
  'Driver=libodbcHDB.so;;ServerNode=<remote_server_name>:
  3<remote_instance_number>15'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

- To enable failover, during creation, add the failover server name to the ServerNode property, separated by a comma. To set session specific connection information, add the sessionVariable property. To enable optimized mode, add the linkeddatabase_mode property.

Sample Code

```
CREATE REMOTE SOURCE <HANA1> ADAPTER hanaodbc
  CONFIGURATION 'DSN=HANA1';ServerNode=<remote_server_name>:
  3<remote_instance_number>15
  [,<failover_server_name>:3<failover_instance_number>15]';
  sessionVariable:<session_variable_name>=?'
  linkeddatabase_mode=optimized'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

- For an existing remote source, to enable failover, set session specific connection information, or enable optimized mode, use ALTER REMOTE SOURCE to add the properties.

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- If you're planing to create the remote source using a DSN entry, it already exists in the ODBC.ini file.

Procedure

1. In the SAP HANA studio, in the *Systems* view of the tenant database, expand the *Provisioning* node within the HANA system.
2. Right-click *Remote Sources* and choose *New Remote Source*.

3. Enter a name for the source. In the *Adapter Name* dropdown list, choose *HANA (ODBC)*.
4. Select your connection Mode.
5. Depending on the mode, mode, enter the required information:

Adapter Properties

Properties	Description
Server	Specifies the server of the remote SAP HANA server. For failover, list the failover server name, separated by a comma. For example, <code>server_name1:30015, failover_server_name1:30015</code> , where 00 is the instance number.
Port	Specifies the port number of the remote SAP HANA server: <code>3<instance_number>15</code> , where <code><instance_number></code> is the instance number of the remote system.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.
Extra Adapter Properties	(Session Connection Information Only) <code>sessionVariable:<session_variable_name>=?</code>

Data Source Name

Properties	Description
To enable failover, session connection information, or optimized mode, add the properties to the DSN section of the <code>ODBC.ini</code> file before creating the remote source.	
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

6. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

- SSO (Kerberos) – All connections to the remote source SAP HANA remote source are authenticated through Kerberos single sign-on (SSO).
7. Choose the *Save this editor* icon in the upper right-hand corner of the screen. Optionally choose the *Test connection* button to verify that the connection to the source was successful.

The data source is now listed under *Remote Sources*. Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

1. In the SAP HANA database explorer, right-click the *Remote Sources* object in your database catalog and click *New Remote Source*.
2. Specify a remote source name.
3. In the *Adapter Name* dropdown list, choose *HANA (ODBC)*.
4. Select your connection mode.
5. Depending on the connection enter the required information:

Adapter Properties

Properties	Description
Server	Specifies the server of the remote SAP HANA server. For failover, list the failover server name, separated by a comma. For example, <code>server_name1:30015, failover_server_name1:30015</code> , where 00 is the instance number.
Port	Specifies the port number of the remote SAP HANA server: <code>3<instance_number>15</code> , where <code><instance_number></code> is the instance number of the remote system.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.
Extra Adapter Properties	(Session Connection Information Only) <code>sessionVariable:<session_variable_name>=?</code>

Data Source Name

Properties	Description
	To enable failover, session connection information, or optimized mode, add the properties to the DSN section of the <code>ODBC.ini</code> file before creating the remote source.
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

6. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

- SSO (Kerberos) – All connections to the remote source SAP HANA remote source are authenticated through Kerberos single sign-on (SSO).
7. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

[Managing Single Sign-On \(SSO\) with Kerberos \[page 1915\]](#)

13.1.5.2 Create an SAP IQ Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to an SAP IQ database.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

An SQL IQ remote source can be used with virtual tables or the linked database feature.

The following syntax examples assume the remote source is configured to use technical user credentials. See *Managing Secondary Credentials* for syntax to use other credential type.

Create a DSN Entry

Context

i Note

These steps are only required if you plan to use a DSN entry to create your remote source.

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (`<sid>adm`), change to the `<sid>adm` \$HOME directory.

2. Create an `.odbc.ini` file if it doesn't already exist.

The `.odbc.ini` must exist, even if empty, regardless of whether a DSN entry is being used.

3. Define one entry in the `.odbc.ini` file for each remote source. For example:

```
[IQ1]
Driver= libdbodbc17_r.so
ServerName=<iq_server_name>
CommLinks=tcpip(host=<iq_machine_name>;port=<IQ_port>)
```

4. Restart the SAP HANA system to apply the changes to the `.INI` file.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- If you're planning to create the remote source using a DSN entry, it already exists in the `ODBC.ini` file.

Procedure

1. In a SQL console, connect to the tenant database.
2. Do one of:
 - If using a DSN entry, execute a `CREATE REMOTE SOURCE` command, referencing the DSN entry in the `.odbc.ini` file.

Sample Code

```
CREATE REMOTE SOURCE MY_IQ1 ADAPTER iqodbc
  CONFIGURATION 'DSN=IQ1'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

- If not using a DSN entry, execute a `CREATE REMOTE SOURCE` command specifying all driver properties.

Sample Code

```
CREATE REMOTE SOURCE MY_IQ2 ADAPTER iqodbc
  CONFIGURATION 'Driver=libdbodbc17_r.so;ServerName=<iq_server_name>
  CommLinks=tcpip(host=<iq_machine_name>;port=<IQ_port>)'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- If you're planning to create the remote source using a DSN entry, it already exists in the `odbc.ini` file.

Procedure

1. In the SAP HANA studio, in the *Systems* view of the tenant database, expand the *Provisioning* node within the HANA system.
2. Right-click *Remote Sources* and choose *New Remote Source*.
3. Enter a name for the source. In the *Adapter Name* dropdown list, choose *IQ (ODBC)*.
4. Select your connection Mode.
5. Depending on the mode, mode, enter the required information:

Adapter Properties

Properties	Description
DML Mode	Specifies if the remote source is readonly (default) or readwrite.
Extra Adapter Properties	Specifies the additional properties to complete the remote connection, in the format: <pre>ServerName=<iq_machine_name>; CommLinks=tcipip(host=<IQ_host>;port=<IQ_port></pre> For example, the additional properties to connect to the demo database would be: <pre>ServerName=<iq_machine_name>_iqdemo; CommLinks=tcipip(host=<iq_machine_name>;port=2638)</pre>

Data Source Name

Properties	Description
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

6. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

7. Choose the *Save this editor* icon in the upper right-hand corner of the screen. Optionally choose the *Test connection* button to verify that the connection to the source was successful.

The data source is now listed under *Remote Sources*. Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

1. In the SAP HANA database explorer, right-click the *Remote Sources* object in your database catalog and click *New Remote Source*.
2. Specify a remote source name.
3. In the *Adapter Name* dropdown list, choose *IQ (ODBC)*.
4. Select your connection mode.
5. Depending on the connection mode, enter the required information:

Adapter Properties

Properties	Description
DML Mode	Specifies if the remote source is readonly (default) or readwrite.
Extra Adapter Properties	Specifies the additional properties to complete the remote connection, in the format: <pre>ServerName=<iq_machine_name>; CommLinks=tcpip(host=<IQ_host>;port=<IQ_port></pre> <p>For example, the additional properties to connect to the demo database would be:</p> <pre>ServerName=<iq_machine_name>_iqdemo; CommLinks=tcpip(host=<iq_machine_name>;port=2638)</pre>

Data Source Name

Properties	Description
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

6. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

7. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

13.1.5.3 Create an SAP Adaptive Server Enterprise Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to an SAP Adaptive Server Enterprise (ASE) database.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

An SAP Adaptive Server Enterprise remote source can be used with virtual tables or the linked database feature. An ASE remote source also supports failover.

The following syntax examples assume the remote source is configured to use technical user credentials. See *Managing Secondary Credentials* for syntax to use other credential type.

Create a DSN Entry

Context

i Note

These steps are only required if you plan to use a DSN entry to create your remote source.

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (<sid>adm), change to the <sid>adm \$HOME directory.
2. Create an .odbc.ini file if it doesn't already exist.

The .odbc.ini must exist, even if empty, regardless of whether a DSN entry is being used.

3. Define one entry in the .odbc.ini file for each remote source. For example:

```
[ASE1]
Server=<ase_machine_name>
Port=<ase_port>
Driver= libsybdrvodb-sqlen8.so
Database=<ase_database_name>
```

4. To enable failover in the DSN entry, add the HASession and AlternativeServers property. For example:

```
[ASE1]
Server=<ase_machine_name>
Port=<ase_port>
Driver= libsybdrvodb-sqlen8.so
Database=<ase_database_name>
HASession=1
AlternateServers=<failover_machine_name>:<failover_port_number>
```

5. Restart the SAP HANA system to apply the changes to the .INI file.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- If you're planning to create the remote source using a DSN entry, it already exists in the ODBC.ini file.

Procedure

1. In a SQL console, connect to the tenant database.
2. Do one of:
 - If using a DSN entry, execute a CREATE REMOTE SOURCE command, referencing the DSN entry in the .odbc.ini file.

Sample Code

```
CREATE REMOTE SOURCE My_ASE1 ADAPTER aseodbc
  CONFIGURATION 'DSN=ASE1'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

- If not using a DSN entry, execute a CREATE REMOTE SOURCE command specifying all driver properties.

Sample Code

```
CREATE REMOTE SOURCE My_ASE2 ADAPTER aseodbc
  CONFIGURATION
  'server=<ase_machine_name>;port=<ase_port>;Driver=libsybdrvodb-
  sqllen8.so;Database=<ase_database_name>'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

- To enable failover, add the HASession=1 and AlternateServers properties to the command.

Sample Code

```
CREATE REMOTE SOURCE My_ASE2 ADAPTER aseodbc
  CONFIGURATION
  'server=<ase_machine_name>;port=<ase_port>;Driver=libsybdrvodb-
  sqllen8.so;Database=<ase_database_name>;

  HASession=1;AlternateServers=<failover_machine_name>:<failover_port_num
  ber>'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

- Use ALTER REMOTE SOURCE to enable failover on an existing ASE remote source.

Sample Code

```
CREATE REMOTE SOURCE My_ASE1 ADAPTER aseodbc
  CONFIGURATION 'DSN=ASE1'
  HASession=1;AlternateServers=<failover_machine_name>:<failover_port_num
  ber>':
```

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- If you're planning to create the remote source using a DSN entry, it already exists in the ODBC.ini file.

Context

If you are using DSN as your connection mode, to enable failover, add the HASession and AlternateServer properties to the DSN section of the ODBC.ini file before creating the remote source.

Procedure

1. In the SAP HANA studio, in the *Systems* view of the tenant database, expand the *Provisioning* node within the HANA system.
2. Right-click *Remote Sources* and choose *New Remote Source*.
3. Enter a name for the source. In the *Adapter Name* dropdown list, choose *ASE (ODBC)*.
4. Select your connection Mode.
5. Depending on the mode, mode, enter the required information:

Adapter Properties

Properties	Description
Server	Specifies the server of the ASE server.
Port	Specifies the port number of the ASE server.
Database Name	Specifies the name of the ASE server.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.
Extra Adapter Properties	(For failover only) Enables automatic failover for the remote source. Enter: <code>HASession=1;AlternateServers=<failover_server>:<failover_port_number></code>

Data Source Name

Properties	Description
To enable failover, add the properties to the DSN section of the <code>ODBC.ini</code> file before creating the remote source.	
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

6. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

7. Choose the *Save this editor* icon in the upper right-hand corner of the screen. Optionally choose the *Test connection* button to verify that the connection to the source was successful.

The data source is now listed under *Remote Sources*. Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

1. In the SAP HANA database explorer, right-click the *Remote Sources* object in your database catalog and click *New Remote Source*.
2. Specify a remote source name.
3. In the *Adapter Name* dropdown list, choose *ASE (ODBC)*.
4. Select your connection mode.
5. Depending on the connection mode, enter the required information:

Adapter Properties

Properties	Description
Server	Specifies the server of the ASE server.
Port	Specifies the port number of the ASE server.
Database Name	Specifies the name of the ASE server.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.
Extra Adapter Properties	(For failover only) Enables automatic failover for the remote source. Enter: <code>HASession=1;AlternateServers=<failover_server>:<failover_port_number></code>

Data Source Name

Properties	Description
To enable failover, add the properties to the DSN section of the <code>ODBC.ini</code> file before creating the remote source.	
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

6. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

7. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

13.1.5.4 Create an SAP Event Stream Processor Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to an Event Stream Processor (ESP) window.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

An SAP Event Stream Processor remote source can be used with virtual tables or the linked database feature.

SAP HANA smart data access does not currently support the BOOLEAN and BINARY data types that exist in ESP. Therefore, any virtual tables created over ESP windows containing these column types either fail or produce incorrect data.

The following syntax examples assume the remote source is configured to use technical user credentials. See *Managing Secondary Credentials* for syntax to use other credential type.

Create a DSN Entry

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (`<sid>adm`), change to the `<sid>adm` \$HOME directory.
2. Create an `.odbc.ini` file if it doesn't already exist.

The `.odbc.ini` must exist, even if empty, regardless of whether a DSN entry is used.

3. Define one entry in the `.odbc.ini` file for each remote source. For example:

```
[ESP1]
```

```
Driver=libstreamingpsqlodbc_lib.so
Database=<esp_workspace_name>/<esp_project_name>
ServerName=<esp_machine_name>
Port=<esp_port_number>
SSLMode=disable
```

4. Restart the SAP HANA system to apply the changes to the .INI file.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- A DSN entry exists in the ODBC.ini file.

Procedure

1. In a SQL console, connect to the tenant database.
2. Do one of:
 - If using a DSN entry, execute a CREATE REMOTE SOURCE command, referencing the DSN entry in the .odbc.ini file.

Sample Code

```
CREATE REMOTE SOURCE MY_ESP1 ADAPTER "odbc"
  CONFIGURATION FILE 'property_esp.ini' CONFIGURATION 'DSN=ESP1'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

- If not using a DSN entry, execute a CREATE REMOTE SOURCE command specifying all driver properties.

Sample Code

```
CREATE REMOTE SOURCE MY_ESP2 ADAPTER "odbc"
  CONFIGURATION FILE 'property_esp.ini'
  CONFIGURATION
  'ServerName=<esp_machine_name>;port=<esp_port_number>;Driver=libstreamin
  glibpsqlodbc_lib.so;
  Database='<esp_workspace_name>/<esp_project_name>;SSLmode=disable'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In the SAP HANA studio, in the *Systems* view of the tenant database, expand the *Provisioning* node within the HANA system.
2. Right-click *Remote Sources* and choose *New Remote Source*.
3. Enter a name for the source. In the *Adapter Name* dropdown list, choose *GENERIC ODBC*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_esp.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Choose the *Save this editor* icon in the upper right-hand corner of the screen. Optionally choose the *Test connection* button to verify that the connection to the source was successful.

The data source is now listed under *Remote Sources*. Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

1. In the SAP HANA database explorer, right-click the *Remote Sources* object in your database catalog and click *New Remote Source*.

2. Specify a remote source name.
3. In the *Adapter Name* dropdown list, choose *GENERIC ODBC*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_esp.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

13.1.5.5 Create an SAP HANA Streaming Analytics Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to a streaming analytics window.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

An SAP HANA streaming analytics remote source can be used with virtual tables or the linked database feature.

SAP HANA smart data access does not currently support the BOOLEAN and BINARY data types that exist in streaming analytics. Therefore, any virtual tables created over streaming analytics windows containing these column types would either fail or produce incorrect data.

The following syntax examples assume the remote source is configured to use technical user credentials. See *Managing Secondary Credentials* for syntax to use other credential type.

Create a DSN Entry

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (<sid>adm), change to the <sid>adm \$HOME directory.
2. Create an `.odbc.ini` file if it doesn't already exist.

The `.odbc.ini` must exist, even if empty, regardless of whether a DSN entry is used.

3. Define one entry in the `.odbc.ini` file for each remote source. For example:

```
[ESP1]
Driver=libstreamingpsqlodbc_lib.so
Database=<esp_workspace_name>/<esp_project_name>
ServerName=<esp_machine_name>
Port=<esp_port_number>
SSLMode=disable
```

4. Restart the SAP HANA system to apply the changes to the `.INI` file.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In a SQL console, connect to the tenant database.
2. Do one of:

- If using a DSN entry, execute a CREATE REMOTE SOURCE command, referencing the DSN entry in the `.odbc.ini` file.

Sample Code

```
CREATE REMOTE SOURCE MY_STREAMING1 ADAPTER "odbc"
  CONFIGURATION FILE 'property_esp.ini' CONFIGURATION
  'DSN=STREAMING1'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

- If not using a DSN entry, execute a CREATE REMOTE SOURCE command specifying all driver properties.

Sample Code

```
CREATE REMOTE SOURCE MY_STREAMING2 ADAPTER "odbc"
  CONFIGURATION FILE 'property_esp.ini'
  CONFIGURATION
  'ServerName=<streaming_machine_name>;port=3<streaming_instance_number>16
  ;
  Driver=libstreaminglibpsqlodbca_lib.so'
  Database='<streaming_workspace_name>/
  <streaming_project_name>;SSLmode=enable'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>;
```

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In the SAP HANA studio, in the *Systems* view of the tenant, expand the *Provisioning* node within the HANA system.
2. Right-click *Remote Sources* and choose *New Remote Source*.
3. Enter a name for the source. In the *Adapter Name* dropdown list, choose *GENERIC ODBC*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_esp.ini</code>

Data Source Name	
Properties	Description
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

- Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

- Choose the [Save this editor](#) icon in the upper right-hand corner of the screen. Optionally choose the [Test connection](#) button to verify that the connection to the source was successful.

The data source is now listed under [Remote Sources](#). Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

- In the SAP HANA database explorer, right-click the [Remote Sources](#) object in your database catalog and click [New Remote Source](#).
- Specify a remote source name.
- In the [Adapter Name](#) dropdown list, choose [GENERIC ODBC](#).
- Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_esp.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

- Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

13.1.5.6 Create an SAP MaxDB Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to an SAP MaxDB database.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

An SAP MaxDB remote source can be used with virtual tables or the linked database feature.

i Note

Enter all passwords in uppercase only, even if it contains lower or mixed case characters. The connection fails when lower or mixed case values are supplied.

The following syntax examples assume the remote source is configured to use technical user credentials. See *Managing Secondary Credentials* for syntax to use other credential type.

Create a DSN Entry

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (<sid>adm), change to the <sid>adm \$HOME directory.
2. Create an .odbc.ini file if it doesn't already exist.

The .odbc.ini must exist, even if empty, regardless of whether a DSN entry is being used.

3. Define one entry in the .odbc.ini file for each remote source. For example:

```
[MaxDB1]
Driver=/opt/MaxDB/lib/libsdboodbcw.so
ServerNode=<maxdb_machine_name>
ServerDB=MAXDB
```

4. Restart the SAP HANA system to apply the changes to the .INI file.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- A DSN entry exists in the ODBC.ini file.

Procedure

1. In a SQL console, connect to the tenant database.
2. Do one of:
 - If using a DSN entry, execute a CREATE REMOTE SOURCE command, referencing the DSN entry in the .odbc.ini file.

Sample Code

```
CREATE REMOTE SOURCE MY_MAXDB1 ADAPTER "odbc"
  CONFIGURATION 'DSN=MaxDB1'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

- If not using a DSN entry, execute a CREATE REMOTE SOURCE command specifying all driver properties.

Sample Code

```
CREATE REMOTE SOURCE MY_MaxDB2 ADAPTER "odbc"
  CONFIGURATION
  'Driver=libsdboodbcw.so;ServerNode=<maxdb_machine_name>;ServerDB=MAXDB'
```

```
with CREDENTIAL TYPE 'PASSWORD' USING
'user=<user_name>;password=<password>';
```

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- A DSN entry exists in the `odbc.ini` file.

Context

Using an SQL console, you can create a remote source using either adapter properties or the data source name. Using SAP HANA studio, only the data source name option is available.

Procedure

1. In the SAP HANA studio, in the *Systems* view of the tenant database, expand the *Provisioning* node within the HANA system.
2. Right-click *Remote Sources* and choose *New Remote Source*.
3. Enter a name for the source. In the *Adapter Name* dropdown list, choose *MaxDB (GENERIC ODBC)*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_maxdb.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Choose the *Save this editor* icon in the upper right-hand corner of the screen. Optionally choose the *Test connection* button to verify that the connection to the source was successful.

The data source is now listed under *Remote Sources*. Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

1. In the SAP HANA database explorer, right-click the *Remote Sources* object in your database catalog and click *New Remote Source*.
2. Specify a remote source name.
3. In the *Adapter Name* dropdown list, choose *MaxDB (GENERIC ODBC)*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_maxdb.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

13.1.5.7 Create a Teradata Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to a Teradata database.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

A Teradata remote source can be used with virtual tables or the linked database feature.

The following syntax examples assume the remote source is configured to use technical user credentials. See *Managing Secondary Credentials* for syntax to use other credential type.

Create a DSN Entry

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (`<sid>adm`), change to the `<sid>adm` \$HOME directory.
2. Create an `.odbc.ini` file if it doesn't already exist.
3. Add the following entries to the `.odbc.ini` file. For example:

```
[ODBC]
InstallDir=<installation_path>/client/ODBC_64
[ODBC Data Sources]
default=tdata.so
TD=tdata.so
```

4. Also in the `.odbc.ini` file, define one entry for each remote source. For example:

```
[TDATA1]
Driver=<installation_path>/client/ODBC_64/lib/tdata.so
DBCName=<server.com>
CharacterSet=UTF8
```

5. Restart the SAP HANA system to apply the changes to the `.INI` file.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In a SQL console, connect to the tenant database.
2. Execute a `CREATE REMOTE SOURCE` command, referencing the DSN entry in the `.odbc.ini` file.

Sample Code

```
CREATE REMOTE SOURCE MY_TDATA1 ADAPTER tdodbc
  CONFIGURATION 'DSN=TDATA1'
  with CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In the SAP HANA studio, in the *Systems* view of the tenant database, expand the *Provisioning* node within the HANA system.
2. Right-click *Remote Sources* and choose *New Remote Source*.
3. Enter a name for the source. In the *Adapter Name* dropdown list, choose *TERADATA (ODBC)*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.

- Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
- Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Choose the *Save this editor* icon in the upper right-hand corner of the screen. Optionally choose the *Test connection* button to verify that the connection to the source was successful.

The data source is now listed under *Remote Sources*. Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

1. In the SAP HANA database explorer, right-click the *Remote Sources* object in your database catalog and click *New Remote Source*.
2. Specify a remote source name.
3. In the *Adapter Name* dropdown list, choose *TERADATA (ODBC)*.
4. Select a connection mode.
5. Depending on your connection mode, enter the required connection information:

Data Source Name	
Properties	Description
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

6. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

7. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

13.1.5.8 Create an SQL Server Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to a SQL Server database.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

An SQL Server remote source can be used with virtual tables or the linked database feature.

The following syntax examples assume the remote source is configured to use technical user credentials. See *Managing Secondary Credentials* for syntax to use other credential type.

Create a DSN Entry

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (`<sid>adm`), change to the `<sid>adm` \$HOME directory.
2. Create an `.odbc.ini` file if it doesn't already exist.
3. Define one entry in the `.odbc.ini` file for each remote source. For example:

```
[MSSQL1]
Driver= /opt/microsoft/msodbcsql/lib64/libmsodbcsql-11.0.so.2260.0
Server=<sql_server_name>,<sql_port>
Database=<sql_database_name>
```

4. Restart the SAP HANA system to apply the changes to the `.INI` file.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In a SQL console, connect to the tenant database.
2. Execute a `CREATE REMOTE SOURCE` command, referencing the DSN entry in the `.odbc.ini` file.

Sample Code

```
CREATE REMOTE SOURCE MY_MSSQL1_ADAPTER odbc
  CONFIGURATION FILE '\property_mss.ini'
  CONFIGURATION 'DSN=MSSQL1'
  with CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In the SAP HANA studio, in the *Systems* view of the tenant database, expand the *Provisioning* node within the HANA system.
2. Right-click *Remote Sources* and choose *New Remote Source*.
3. Enter a name for the source. In the *Adapter Name* dropdown list, choose *MSSQL (GENERIC (ODBC))*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_mss.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.

Data Source Name	
Properties	Description
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Choose the *Save this editor* icon in the upper right-hand corner of the screen. Optionally choose the *Test connection* button to verify that the connection to the source was successful.

The data source is now listed under *Remote Sources*. Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

1. In the SAP HANA database explorer, right-click the *Remote Sources* object in your database catalog and click *New Remote Source*.
2. Specify a remote source name.
3. In the *Adapter Name* dropdown list, choose *MSSQL (GENERIC (ODBC))*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_mss.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

13.1.5.9 Create an IBM DB2 Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to a DB2 database.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

An IBM DB2 remote source can be used with virtual tables or the linked database feature.

The following syntax examples assume the remote source is configured to use technical user credentials. See *Managing Secondary Credentials* for syntax to use other credential type.

Create a DSN Entry

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (`<sid>adm`), change to the `<sid>adm` \$HOME directory.
2. Create an `.odbc.ini` file if it doesn't already exist.
3. Define a DSN entry in `.odbc.ini` file for each remote source. For example:

```
[DB2_1]
Driver=/opt/ibm/db2/odbc_cli/clidriver/lib/libdb2o.so.1
Description=TEST_DB2
```

- Restart the SAP HANA system to apply the changes to the .INI file.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- A DSN entry exists in the ODBC.ini file.

Procedure

- In a SQL console, connect to the tenant database.
- Execute a CREATE REMOTE SOURCE command, referencing the DSN entry in the .odbc.ini file.

Sample Code

```
CREATE REMOTE SOURCE MY_DB2_11 ADAPTER odbc
  CONFIGURATION FILE 'property_db2.ini' CONFIGURATION 'DSN=DB2_1'
  WITH CREDENTIAL TYPE 'PASSWORD' USING
  'user=<user_name>;password=<password>';
```

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- A DSN entry exists in the ODBC.ini file.

Procedure

- In the SAP HANA studio, in the *Systems* view of the tenant database, expand the *Provisioning* node within the HANA system.
- Right-click *Remote Sources* and choose *New Remote Source*.
- Enter a name for the source. In the *Adapter Name* dropdown list, choose *DB2 (GENERIC ODBC)*.
- Enter the required connection information:

Data Source Name

Properties	Description
Configuration File	Specify <code>property_db2.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Choose the *Save this editor* icon in the upper right-hand corner of the screen. Optionally choose the *Test connection* button to verify that the connection to the source was successful.

The data source is now listed under *Remote Sources*. Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

1. In the SAP HANA database explorer, right-click the *Remote Sources* object in your database catalog and click *New Remote Source*.
2. Specify a remote source name.
3. In the *Adapter Name* dropdown list, choose *DB2 (GENERIC ODBC)*.
4. Enter the required connection information:

Data Source Name

Properties	Description
Configuration File	Specify <code>property_db2.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

13.1.5.10 Create an Oracle Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to an Oracle database.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

An Oracle remote source can be used with virtual tables or the linked database feature.

Oracle remote sources support failover. Refer to your Oracle product documentation for information on configuring failover.

Oracle remote sources do not support empty strings. Values inserted into a virtual table that are generated from an Oracle remote source are transformed into NULL values if they are empty strings. This behavior also impacts some of the smart data access optimization techniques (for example, join relocation).

The following syntax examples assume the remote source is configured to use technical user credentials. See *Managing Secondary Credentials* for syntax to use other credential type.

Create a DSN Entry

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (<sid>adm), change to the <sid>adm \$HOME directory.
2. In the <sid>adm \$HOME folder, create a `tnsnames.ora` file if it doesn't already exist. Define one entry for each remote source. For example:

```
ORCL1=
(DESCRIPTION =
(AADDRESS = (PROTOCOL = TCP) (HOST = <oracle_hostname>) (PORT = 1521) )
(CONNECT_DATA = (SERVER = DEDICATED) (SERVICE_NAME =
<oracle_machine_name>) ) )
```

3. Also in the <sid>adm \$HOME folder, create an `.odbc.ini` file if it doesn't already exist and define one entry for each remote source. For example:

```
[ORA12C_1]
Driver=<install_oracle_driver_folder>/instantclient_12_1/libsqora.so.12.1
ServerName=ORCL
```

4. To enable failover, add the `HASession` and `AlternateServers` properties to the command. For example:

```
[ORA12C_1]
Driver=<install_oracle_driver_folder>/instantclient_12_1/libsqora.so.12.1
ServerName=ORCL
HASession=1;AlternateServers=<failover_machine_name>:<failover_port_number>'
```

5. Restart the SAP HANA system to apply the changes to the files.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In a SQL console, connect to the tenant database.
2. Do one of:
 - Execute the `CREATE REMOTE SOURCE` command, referencing the DSN entry in the `.odbc.ini` file.

Sample Code

```
CREATE REMOTE SOURCE MY_ORA12C_1 ADAPTER "odbc"
CONFIGURATION FILE 'property_orcl.ini' CONFIGURATION 'DSN=ORA12C_1'
```

```
WITH CREDENTIAL TYPE 'PASSWORD' USING  
'user=<user_name>;password=<password>';
```

- To enable failover, add the `HASession=1` and `AlternateServers` properties to the command.

Sample Code

```
CREATE REMOTE SOURCE My_ORA12C_1 ADAPTER "odbc"  
  CONFIGURATION FILE 'property_orcl.ini' CONFIGURATION 'DSN=ORA12C_1'  
  
HASession=1;AlternateServers=<failover_machine_name>:<failover_port_number>  
'  
  WITH CREDENTIAL TYPE 'PASSWORD' USING  
'user=<user_name>;password=<password>';
```

- Use `ALTER REMOTE SOURCE` to enable failover on an existing Oracle remote source.

Sample Code

```
ALTER REMOTE SOURCE My_ORA12C_1 ADAPTER "odbc"  
  CONFIGURATION FILE 'property_orcl.ini' CONFIGURATION  
'DSN=ORA12C_1'; HASession=1;  
  AlternateServers=<failover_machine_name>:<failover_port_number>'
```

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Context

To enable failover, add the `HASession` and `AlternateServer` properties to the DSN section of the `ODBC.ini` file before creating the remote source.

Procedure

1. In the SAP HANA studio, in the *Systems* view of the tenant database, expand the *Provisioning* node within the HANA system.
2. Right-click *Remote Sources* and choose *New Remote Source*.
3. Enter a name for the source. In the *Adapter Name* dropdown list, choose *ORACLE (GENERIC ODBC)*.
4. Enter the required connection information:

Data Source Name

Properties	Description
	To enable failover, add the properties to the DSN section of the <code>ODBC.ini</code> file before creating the remote source.
Configuration File	Specify <code>property_orcl.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Choose the [Save this editor](#) icon in the upper right-hand corner of the screen. Optionally choose the [Test connection](#) button to verify that the connection to the source was successful.

The data source is now listed under [Remote Sources](#). Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

1. In the SAP HANA database explorer, right-click the [Remote Sources](#) object in your database catalog and click [New Remote Source](#).
2. Specify a remote source name.
3. In the [Adapter Name](#) dropdown list, choose [ORACLE \(GENERIC ODBC\)](#).
4. Enter the required connection information:

Data Source Name

Properties	Description
	To enable failover, add the properties to the DSN section of the <code>ODBC.ini</code> file before creating the remote source.
Configuration File	Specify <code>property_orcl.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.

- Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

13.1.5.11 Create an IBM Netezza Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to an IBM Netezza database.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

An IBM Netezza remote source can be used with virtual tables or the linked database feature.

The following syntax examples assume the remote source is configured to use technical user credentials. See *Managing Secondary Credentials* for syntax to use other credential type.

Create a DSN Entry

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (`<sid>adm`), change to the `<sid>adm` \$HOME directory.

2. Add the following entry to the `odbcinst.ini` file:

```
[ODBC Drivers]
NetezzaSQL = Installed
[NetezzaSQL]
Driver          = /usr/local/nz/lib64/libnzodbc.so
Setup          = /usr/local/nz/lib64/libnzodbc.so
APILevel       = 1
ConnectFunctions = YYN
Description    = Netezza ODBC driver
DriverODBCVer  = 03.51
DebugLogging   = false
LogPath        = /tmp
UnicodeTranslationOption = utf16
CharacterTranslationOption = all
PreFetch       = 25600
Socket         = 16384
```

i Note

Set the `PreFetch` property to a large value since IBM Netezza does not support multicursor on same connection.

3. Create an `.odbc.ini` file if it doesn't already exist.
4. Define a DSN entry in `.odbc.ini` file for each remote source. For example:

```
[NTZ1]
Driver = /usr/local/nz/lib64/libnzodbc.so
Servername = <ibm_machine_name>
Port = 5480
database = <ibm_database_name>
Username = <username>
Password = <password>
UnicodeTranslationOption = utf16
CharacterTranslationOption = all
```

i Note

The `UnicodeTranslationOption` entry is required for HANA to connect successfully to an IBM Netezza remote source.

5. Restart the SAP HANA system to apply the changes to the `.INI` file.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In a SQL console, connect to the tenant database.
2. Execute a CREATE REMOTE SOURCE command, referencing the DSN entry in the `.odbc.ini` file. For example:

Sample Code

```
CREATE REMOTE SOURCE MY_NTZ1 ADAPTER "odbc"  
    CONFIGURATION FILE 'property_ntz.ini' CONFIGURATION 'DSN=NTZ1'  
    WITH CREDENTIAL TYPE 'PASSWORD' USING 'user=<username>;password=  
<password>';
```

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In the SAP HANA studio, in the *Systems* view of the tenant database, expand the *Provisioning* node within the HANA system.
2. Right-click *Remote Sources* and choose *New Remote Source*.
3. Enter a name for the source. In the *Adapter Name* dropdown list, choose *NETEZZA (GENERIC ODBC)*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_ntz.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Choose the *Save this editor* icon in the upper right-hand corner of the screen. Optionally choose the *Test connection* button to verify that the connection to the source was successful.

The data source is now listed under *Remote Sources*. Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

1. In the SAP HANA database explorer, right-click the *Remote Sources* object in your database catalog and click *New Remote Source*.
2. Specify a remote source name.
3. In the *Adapter Name* dropdown list, choose *NETEZZA (GENERIC ODBC)*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_ntz.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

i Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

13.1.5.12 Create a Google BigQuery Remote Source

Use SAP HANA studio, SAP HANA database explorer, or SQL syntax to create a remote source to a Google BigQuery database.

Prerequisites

- You have the CREATE REMOTE SOURCE system privilege.
- If you're planning to create the remote source using a DSN entry, it must already exist in the `ODBC.ini` file.
- The remote data source is reachable by the network from the computer you are using.

Context

A Google BigQuery remote source can be used with virtual tables or the linked database feature.

Defining credentials when creating a Google BigQuery remote source using SAP HANA Studio or SAP HANA Database Explorer are not supported.

Procedure

1. Enable safe mode for Google BigQuery.
 - a. If the instance does not yet have a scriptserver, connect to SYSTEMDB and in a console window, execute:

```
ALTER DATABASE <tenant_db_name> ADD 'scriptserver'.
```

- b. Connect to the tenant database. Change the value of `odbc_adapters_in_scriptserver` parameter in `indexserver.ini`.

```
ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini', 'SYSTEM')  
    SET ('smart_data_access', 'odbc_adapters_in_scriptserver') =  
    'bigquery';
```

2. Change to the `/usr/sap/<sid>/HDB<instance_number>/exe/config` folder.
3. Open the `property_bq.ini` file and ensure the Make sure that the `PROP_USE_UNIX_DRIVER_MANAGER` property is set to `true`. This forces SAP HANA to use unixODBC driver manager.

```
PROP_USE_UNIX_DRIVER_MANAGER : true
```

Create a DSN Entry

Procedure

1. Log on to the SAP HANA host as the SAP HANA software owner (<sid>adm), change to the <sid>adm \$HOME directory.
2. Create an .odbc.ini file if it doesn't already exist and define one entry for each remote source. Use this template as a guideline.

```
[GoogleBQ]
# Description: DSN Description.
# This key is not necessary and is only to give a description of the data
source.
Description=Simba ODBC Driver for Google BigQuery (64-bit) DSN

# Driver: The location where the ODBC driver is installed to.
Driver=/opt/simba/googlebigqueryodbc/lib/64/libgooglebigqueryodbc_sb64.so

# These values can be set here, or on the connection string.
# Catalog: The catalog to connect to. This is a required setting.
Catalog=

# SQLDialect: The SQL Dialect to use. There are two SQL dialects:
# 0 = BigQuery Legacy SQL
# 1 = BigQuery Standard SQL (SQL 11)
SQLDialect=1

# OAuth Mechanism: The OAuth mechanism to use. There are two choices:
# 0 = Service Authentication
# 1 = User Authentication
#
# This is a required setting.
OAuthMechanism=0

# RefreshToken: The Refresh Token used. This can be generated from the
Windows connection dialog.
# It can also be generated by executing the following steps:
# 1. Get an Authentication by logging into Google from the following URL:
# https://accounts.google.com/o/oauth2/auth?scope=https://www.googleapis.com/
auth/bigquery&response_type=
# code&redirect_uri=urn:ietf:wg:oauth:
2.0:oob&client_id=977385342095.apps.googleusercontent.com&hl=en
# &from_login=1&as=76356ac9e8ce640b&pli=1&authuser=0
# 2. Run the get_refresh_token.sh shell script and pass in the Authentication
Token received in step 1.
# 3. Copy the Refresh Token (the text on the right-side of the colon, without
the trailing or leading spaces)
# from the output of the script. This is a required setting.
#RefreshToken=

# Email: For Service Authentication, this is a required setting. It is your
GENERATED service account email
# (not a typical Gmail account).
# It is unique and associated with at least one public/private key pair.
Email=

# KeyFile Path: For Service Authentication, this is a required setting. This
is the path to the stored keyfile (.p12).
KeyFilePath=

# Used to specify the full path of the PEM formatted file containing trusted
SSL CA certificates.
```

```

# If an empty string is passed in for the configuration, the driver expects
the trusted SSL CA
# certificates can be found in the file named cacerts.pem located in the same
directory as the
# driver's shared library.
#TrustedCerts=

# AllowLargeResults: When set to 1, the driver allows for result sets in
responses to be larger than 128 MB.
AllowLargeResults=0

# LargeResultsDataSetId: DataSetId to store temporary tables created. This
is a required setting if
# AllowLargeResults is set to 1.
LargeResultsDataSetId=_bqodbc_temp_tables

# LargeResultsTempTableExpirationTime: Time in milliseconds before the
temporary tables created expire.
# This is a required setting if AllowLargeResults is set to 1.
LargeResultsTempTableExpirationTime=3600000

```

3. Change to the `/usr/sap/<sid>/HDB<instance_number>/exe/config` folder.
4. Open the `property_bq.ini` file and ensure the Make sure that the `PROP_USE_UNIX_DRIVER_MANAGER` property is set to `true`. This forces SAP HANA to use unixODBC driver manager.

```
PROP_USE_UNIX_DRIVER_MANAGER : true
```

5. Restart the SAP HANA system to apply the changes to the `.INI` file.

Create a Remote Data Source Using SQL Syntax

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In a SQL console, connect to the tenant database.
2. Execute a `CREATE REMOTE SOURCE` command, referencing the DSN entry in the `.odbc.ini` file. For example:

Sample Code

```
CREATE REMOTE SOURCE BigQ ADAPTER "odbc"
CONFIGURATION FILE 'property_bq.ini' CONFIGURATION 'DSN=GoogleBQ'
```

Create a Remote Data Source Using SAP HANA Studio

Prerequisites

- A DSN entry exists in the `ODBC.ini` file.

Procedure

1. In the SAP HANA studio, in the *Systems* view of the tenant database, expand the *Provisioning* node within the HANA system.
2. Right-click *Remote Sources* and choose *New Remote Source*.
3. Enter a name for the source. In the *Adapter Name* dropdown list, choose *BIGQUERY (GENERIC ODBC)*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_bq.ini</code>
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Choose the *Save this editor* icon in the upper right-hand corner of the screen. Optionally choose the *Test connection* button to verify that the connection to the source was successful.

The data source is now listed under *Remote Sources*. Expand the data source to see the users and tables.

Create a Remote Source Using SAP HANA Database Explorer

Procedure

1. In the SAP HANA database explorer, right-click the *Remote Sources* object in your database catalog and click *New Remote Source*.
2. Specify a remote source name.
3. In the *Adapter Name* dropdown list, choose *BIGQUERY (GENERIC ODBC)*.
4. Enter the required connection information:

Data Source Name	
Properties	Description
Configuration File	Specify <code>property_bq.ini</code>

Data Source Name

Properties	Description
Data Source Name	Specifies the DSN as defined in the <code>odbc.ini</code> file.
DML Mode	Specifies if the remote source is readonly (default) or readwrite.

5. Specify the user credentials to connect to the remote source.
 - Technical user – All connections to the remote data source share the same credential for the data source. Specify a valid SAP HANA user and password to connect to the remote source.
 - Secondary credentials - One credential per user per data source.

Note

At least one secondary credential should exist before creating the remote source. If no secondary credentials exist, credential mode is set to None. Once a secondary credential is created, credential mode switches to secondary credentials.

6. Click *OK*.

Related Information

[Managing Secondary Credentials \[page 1912\]](#)

13.1.6 Modifying a Remote Source

Modify an existing remote source.

Prerequisites

One of the following:

- You created the remote source.
- You have the CREATE REMOTE SOURCE privilege.

Modify a Remote Source Using a SQL Console

Procedure

In a SQL console, execute:

```
ALTER REMOTE SOURCE <remote_source_name> <adapter_clause>
```

```
[ <credential_clause> ]
```

❖ Example

This example changes the port for the remote source `remote1` to a HANA remote source from 30115 to 30315.

≡ Sample Code

```
CREATE REMOTE SOURCE remote1 ADAPTER "hanaodbc"  
  CONFIGURATION 'Driver=libodbcHDB.so;ServerNode=my_machine:30115'  
  WITH CREDENTIAL TYPE 'PASSWORD' USING  
  'user=<user_name>;password=<password>';  
  
ALTER REMOTE SOURCE remote1 ADAPTER "hanaodbc"  
  CONFIGURATION 'Driver=libodbcHDB.so;ServerNode=<machine_name>:30315'
```

Modify a Remote Source Using SAP HANA Studio

Procedure

1. In the *Systems* view, expand the **Provisioning** > *Remote Sources* node.
2. Select the remote source to modify and choose *Open Definition* from the context menu.
3. Make the changes and click the *Save the Editor* icon (💾).
4. Click the *Test connection* icon (🟢) to verify the connection is still valid.

13.1.7 Dropping a Remote Source

Remove an existing remote source.

Prerequisites

One of the following:

- You created the remote source.
- You have the CREATE REMOTE SOURCE system privilege.

Drop a Remote Source Using a SQL Console

Procedure

In a SQL window, execute:

```
DROP REMOTE SOURCE <remote_source_name> CASCADE;
```

Drop a Remote Source Using SAP HANA Studio

Procedure

1. In the *Systems* view, expand the **Provisioning > Remote Sources** node.
2. Select the remote source to delete and choose *Delete* from the context menu.
3. Choose *Yes* to confirm.

Results

The remote source disappears from the *Systems* view, including all dependent virtual tables.

Drop a Remote Source Using SAP HANA Database Explorer

Procedure

In the catalog browser item list, right-click the remote source you want to drop, and click *Delete*.

13.1.8 Listing Remote Sources

Provides a list of remote sources you have privilege to.

Procedure

1. In a SQL console, connect to the tenant database.

2. Execute:

```
select * from "SYS"."REMOTE_SOURCE"
```

13.1.9 Managing Secondary Credentials

Secondary credentials let you assign different credentials to different users when using a remote source.

To access a remote database requires valid credentials on the database. All actions performed on the remote database are executed using these credentials and the privileges associated with them.

When creating a remote source, you can define a remote user name and password, called the technical user. With this configuration, all users using the remote source use the same technical user credentials to access to the remote database. If a user with secondary credentials accesses a remote source with a technical user defined, the technical user credentials are used and the secondary credentials are ignored.

If you want to associate different remote credentials with individual users, configure the remote source to use secondary credentials. With this configuration, users without secondary credentials can't access the remote source.

Dropping a remote source automatically drops all credentials, including secondary credentials associated with the remote source.

No privileges are required to manage your own credentials, but the CREDENTIAL ADMIN privilege is required to manage other credentials. Management of technical user credentials can be done by the owner of the remote source, or any user with the CREATE REMOTE SOURCE or CREDENTIAL ADMIN privilege.

13.1.9.1 Create a Remote Source Using Secondary Credentials

Create a remote source that uses secondary credentials to control access at the user level.

Prerequisites

- Requires CREATE REMOTE SOURCE system privilege to manage remote sources.
- Requires CREDENTIAL ADMIN privilege to manage other credentials. No privileges are required to manage your own credentials.

Context

Though not mandatory, creating the secondary credentials before creating the remote source speeds the process. Secondary credentials can only be created using SQL.

Create a Remote Source Using a SQL Console

Procedure

1. Create secondary credentials for each user to access the remote source, including user SYSTEM.

```
CREATE CREDENTIAL FOR USER <HANA_user> COMPONENT 'SAPHANAFEDERATION' PURPOSE
'<remote__source_name>' TYPE 'PASSWORD' USING
'user=<remote_user_name>;password=<remote_user_password>';
```

2. Create a remote source without credentials. The CONFIGURATION syntax to create a remote source depends upon the connection mode and ODBC driver. See *Creating a Remote Source* for the required syntax for each driver.

```
CREATE REMOTE SOURCE <remote_source_name> ADAPTER <ODDBC_driver>
CONFIGURATION <driver_specific_syntax>;
```

❁ Example

Create secondary credentials on remote source HANA1 for user user1:

```
CREATE CREDENTIAL FOR USER user1 COMPONENT 'SAPHANAFEDERATION' PURPOSE 'HANA1'
TYPE 'PASSWORD' USING 'user=<remote_user>;password=<remote_password>';
```

Create remote source HANA1 to SAP HANA system HA1, instance 00, without credentials:

```
CREATE REMOTE SOURCE HANA1 ADAPTER hanaodbc
CONFIGURATION 'Driver=libodbcHDB.so;ServerNode=<machine_name>;30015';
```

Create a Remote Source Using SAP HANA Studio

Procedure

1. In a SQL console, create secondary credentials for each user to access the remote source, including user SYSTEM.

```
CREATE CREDENTIAL FOR USER <HANA_user> COMPONENT 'SAPHANAFEDERATION' PURPOSE
'<remote__source_name>' TYPE 'PASSWORD' USING
'user=<remote_user_name>;password=<remote_user_password>';
```

2. In the *Systems* view, expand the *Provisioning* node.
3. Select *Remote Sources* and choose *New Remote Source* from the context menu.
4. Enter a name for the source. In the *Adapter Name* dropdown list, choose applicable ODBC driver.
5. Enter the properties for the specified ODBC driver. See *Creating a Remote Source* for driver-specific properties. For Credential mode, choose *Secondary credentials*.
6. Choose the *Save this editor* icon in the upper right-hand corner of the screen.

If no secondary credentials exist for the remote source, a message appears indicating that only NONE is valid for the remote source. Once the first secondary is created, the mode automatically changes to secondary credentials.

Results

The data source is now listed under *Remote Sources*. Expanding the data source displays the users and tables.

i Note

If the logged on user doesn't have secondary credentials on the remote source, an error appears indicating credentials cannot be found.

Related Information

[Creating a Remote Source \[page 1863\]](#)

13.1.9.2 Convert an Existing Remote Source to Use Secondary Credentials

Change the credentials mode of an existing remote source to use secondary credentials.

Prerequisites

- Requires CREDENTIAL ADMIN privilege to manage other credentials. No privileges are required to manage your own credentials.
- Requires CREATE REMOTE SOURCE or DATA ADMIN privilege to manage remote sources.

Convert A Remote Source Using SQL Console

Procedure

1. Create secondary credentials for each user to access the remote source, including user SYSTEM.

```
CREATE CREDENTIAL FOR USER <HANA_user> COMPONENT 'SAPHANAFEDERATION' PURPOSE
'<remote_source_name>' TYPE 'PASSWORD' USING
'user=<remote_user_name>;password=<remote_user_password>';
```

2. Execute:

```
ALTER REMOTE SOURCE <remote__source_name> DROP CREDENTIAL TYPE 'PASSWORD';
```

Convert A Remote Source Using SAP HANA Studio

Procedure

1. In a SQL console, create secondary credentials for each user to access the remote source, including user SYSTEM.

```
CREATE CREDENTIAL FOR USER <HANA_user> COMPONENT 'SAPHANAFEDERATION' PURPOSE  
'<remote__source_name>' TYPE 'PASSWORD' USING  
'user=<remote_user_name>;password=<remote_user_password>';
```

2. In the *Systems* view, expand the **Provisioning > Remote Sources** node.
3. Select the remote source to convert choose *Open Definition* from the context menu.
4. In the *Credentials Mode* dropdown list, choose *Secondary credentials*.
5. Choose the *Save this editor* icon in the upper right-hand corner of the screen.

13.1.10 Managing Single Sign-On (SSO) with Kerberos

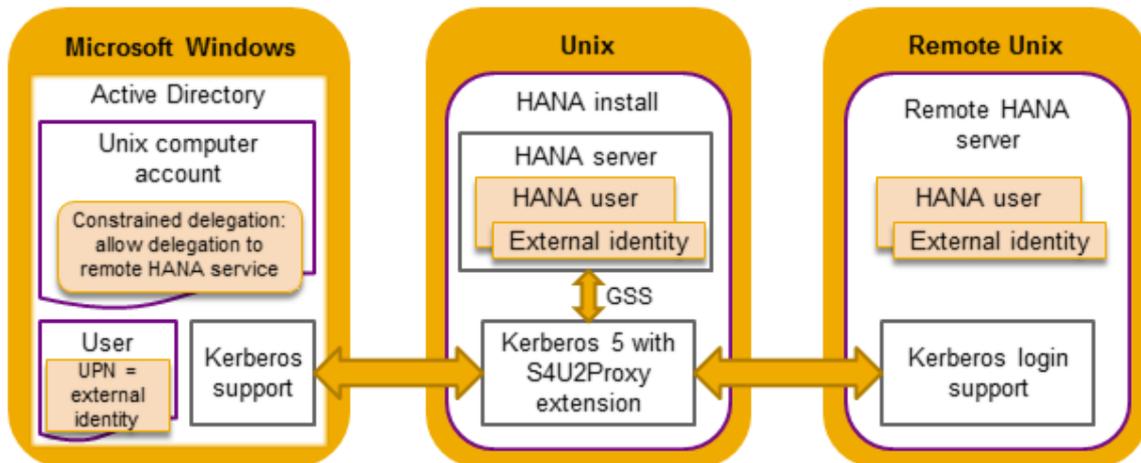
SAP HANA smart data access supports single sign-on with Kerberos for connections to SAP HANA remote sources. Using Kerberos constrained delegation and protocol transition allows SAP HANA users to be authenticated automatically on Microsoft Windows Active Directory, without having to provide a password (SSO mode).

Prerequisites

Microsoft Windows Server, version 2003 or later.

Architecture Overview

The Kerberos platform architecture used in SSO authentication for connections to SAP HANA remote sources is shown below. Protocol transition is assured by Kerberos 5's S4U2Proxy extension:



1. The source SAP HANA server is authenticated only once as a computer using the `krb5_host.keytab`.
2. Users sign on to the source SAP HANA server using an authentication protocol.
3. The source SAP HANA server requests a constrained delegation ticket in its name for the SAP HANA user external identity.
4. The connection to the remote SAP HANA server is authenticated with the constrained delegation ticket.
5. In the authentication validation process, the constrained delegation ticket is validated against the remote SAP HANA service SPN.

Kerberos 5 is installed automatically together with SAP HANA. It contains the S4U (Service for User) extension needed for user impersonation and constrained delegation. Constrained delegation means that delegation can be done only to a predefined set of services. For the purposes of protocol transition, the computer on which the server is installed needs to be entrusted by the Microsoft Windows Active Directory for delegation.

Note that the Kerberos platform is used in SAP HANA for authentication only and not for session management.

Configuration

The Kerberos configuration is defined in the following configuration files:

Configuration File	Description
<code><sidadm home>/etc/krb5_hdb.conf</code>	Configuration of the Kerberos realm to be used with the SAP HANA server installed under <code><SID>adm</code> .

Configuration File	Description
<sidadm home>/etc/krb5_hdb.keytab	List of service keys required to authenticate the services on the Kerberos server.
<sidadm home>/etc/krb5_host.keytab	One entry only to authenticate the host on the Kerberos server for the purpose of delegation.

If the files are present in the <sidadm home>/etc folder, the configuration is automatically taken from there, otherwise the default OS configuration in /etc/krb5.conf and /etc/krb5.keytab is used instead.

For a custom setup of Kerberos, you can overwrite the following variables in /usr/sap/<SID>/home/.customer.sh: KRB5_CONFIG, KRB5_KTNAME, KRB5_CLIENT_KTNAME. For example:

Sample Code

```
export KRB5_CONFIG=<conf file>
export KRB5_KTNAME=<hdb keytab file>
export KRB5_CLIENT_KTNAME=<host keytab file>
```

13.1.10.1 Configure Kerberos On SAP HANA Source

On the source SAP HANA server, configure Kerberos to support constrained delegation.

Procedure

1. Configure the Kerberos realm to be used with the SAP HANA server and enable delegation by setting the `forwardable` parameter for Kerberos service tickets to `true` in the `krb5_hdb.conf` file.
2. On the Microsoft Windows Active Directory server, create a Windows Domain account for the SAP HANA server computer and map a host service principal name (SPN) to it.
3. Add the `hdb`.
4. Add a keytab entry for the `hdb` service. The keytab stores the keys needed by the SAP HANA server to take part in the authentication protocol. service of a remote SAP HANA server to a Microsoft Windows Active Directory account in order to be able to log in to the remote SAP HANA server using Kerberos. Enable constrained delegation and protocol transition for your remote SAP HANA server in the Active Directory Users and Computers application.

Results

For more information about how to set up SSO for SAP HANA smart data access using Kerberos and Microsoft Windows Active Directory, see the **SAP HANA Smart Data Access Single Sign-On Guide** attached to SAP Note [2303807](#).

13.1.10.2 Create a Remote Data Source Using Kerberos Authentication

Connect users to an SAP HANA remote source using single sign-on (SSO) with Kerberos.

Context

Kerberos authentication is primarily used as an alternative to the technical user credential type. You can declare it as either a global credential type for the remote source or as the individual type for a given user. If a user with user level credentials defined and the remote source has global credentials defined, the global credentials are used; the user level credentials are ignored on the remote source.

A user can modify their own credentials to use Kerberos, but require the CREDENTIAL ADMIN privilege to modify other's credentials.

Procedure

1. Do one of:

To:	Execute:
Create global credentials	<pre>CREATE CREDENTIAL FOR COMPONENT 'SAPHANAFEDERATION' PURPOSE <remote_source_name> TYPE 'KERBEROS';</pre>
Create user level credentials	<pre>CREATE CREDENTIAL FOR USER <user_name> COMPONENT 'SAPHANAFEDERATION' PURPOSE <remote_source_name> TYPE 'KERBEROS';</pre>

2. Create a remote data source using credential type KERBEROS. See *Creating a Remote Source* for steps for each ODBC driver.

Related Information

[Creating a Remote Source \[page 1863\]](#)

13.1.11 Enabling Read-Write Access to a Remote Source

The DML mode property specifies whether read-only or read-write access to the remote source is allowed. A remote source is by default read-only.

Prerequisites

- Requires the CREATE REMOTE SOURCE system privilege.

Context

You can set the property in the SAP HANA Studio remote source editor or using a DDL command. When the property DML_MODE is set to READONLY on a remote source, INSERT, UPDATE, and DELETE operations cannot be executed on virtual tables created on this remote source. Note that the read-only option has a positive impact on performance for SELECT queries.

Enable Read-write Access Using SQL Console

Procedure

1. To display the current value of the DML mode option, execute:

```
SELECT SUBSTR_AFTER (CONNECTION_INFO, 'dml_mode">' ) "DML Mode"  
  from PUBLIC.REMOTE_SOURCES  
  where REMOTE_SOURCE_NAME= '<remote_source_name>';
```

2. To modify the DML mode option, in a console window, execute:

```
CREATE REMOTE SOURCE <remote_source_name> ADAPTER "<adapter_name>"  
  CONFIGURATION 'ServerNode=<host_name>:<port>;Driver=<library_name>;  
  DML_MODE=READONLY' WITH CREDENTIAL TYPE 'PASSWORD'  
  USING 'user=<user_name>;password=<password>'
```

Enable Read-write Access Using Hana Studio

Procedure

1. In the *Systems* view, expand the  *Provisioning*  *Remote Sources* node.
2. Right-click the remote source and select *Open Definition*.

3. Click the right edge of the *DML Mode* field to enable the dropdown arrow.
4. Left-click the arrow and select the new value from the list.
5. Save the change.

13.1.12 Managing Virtual Tables

Use SAP HANA Data Explorer, SAP HANA Studio, or SQL syntax to manage virtual tables.

i Note

- Managing virtual tables is not applicable when using linked database.
- Some management tasks are not available in both management tools.

13.1.12.1 Managing Virtual Tables Using SAP HANA Studio

Create, view table content and definition, and delete virtual tables using SAP HANA studio.

i Note

Refreshing the contents of a remote table is only available through SQL syntax. See *Refresh Virtual Tables*.

Related Information

[Refresh Virtual Tables \[page 1925\]](#)

13.1.12.1.1 Create Virtual Tables by Remote Object

Add a virtual table from the remote object in the remote source SAP HANA studio.

Prerequisites

One of the following:

- You created the remote source.
- You have the CREATE VIRTUAL TABLE object privilege on the remote source created by another user.

Procedure

1. In the *Systems* view of the local system, expand the remote source to display the schema containing the remote table to be used: ► *Provisioning* ► *Remote Sources* ► *<remote-source-name>* ► *<database_name>* ► *<schema>* ►.

A list of tables on the remote source for the selected schema appears.

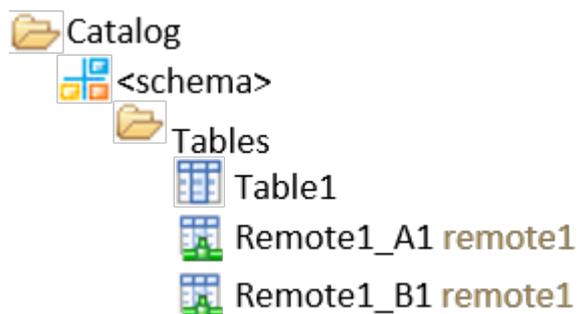
2. Right-click the table to create the virtual table from and choose *Add as Virtual Table* from the context menu.

The *Create Virtual Table* dialog box appears. The table name and schema are automatically filled in. The default virtual table name is *<remote-source-name>_<remote-object-name>*, but you can override the name, if needed. The schema is that of the logged on user.

3. Choose *Create*.
A confirmation message appears when the virtual table is successfully created.

Results

The new virtual table appears in the *Systems* view of the local system, under ► *Catalog* ► *<schema>* ► *Tables* ►.



The list includes both local and virtual tables, indicated by the  icon.

13.1.12.1.2 Create a Virtual Table by Schema

Add a virtual table from the catalog of the local schema using SAP HANA studio.

Prerequisites

One of the following:

- You created the remote source.
- You have the CREATE VIRTUAL TABLE object privilege on the remote source created by another user.

Procedure

1. In the *Systems* view of the local system, expand the catalog to display the schema for the virtual table being created: **Catalog** > <schema> > >.

2. Right-click *Tables* and choose *New Virtual Table* from the context menu.

The *New Virtual Table* tab appears.

3. Select *Browse*.

4. Expand the remote source and <database>.

5. Expand the <schema_name> containing the remote table.

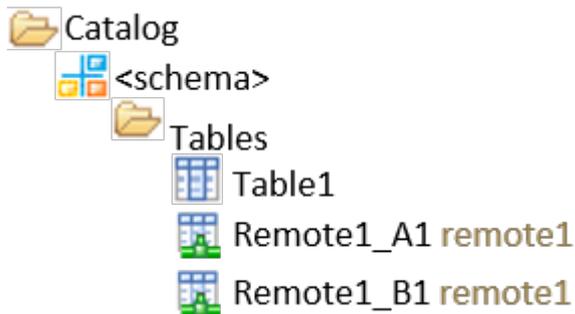
6. Choose the remote table and click *OK*.

The virtual table name appears in the *Table name* field using the default naming convention <remote-source-name>_<remote-table-name>. Override the name, if needed. The column definitions of the table also appear, but cannot be modified.

7. Choose the *Save the Editor* icon () in the upper right corner of the screen to create the virtual table.

Results

When complete, the new virtual table appears in the *Tables* folder of the specified schema.



13.1.12.1.3 List Virtual Tables By Schema

Display the virtual tables of a remote source by schema using SAP HANA studio.

Context

You cannot change the definition or contents while viewing.

Procedure

In the *Systems* view of the local system, expand the folders ► *Catalog* ► *<schema_name>* ► *Tables* ►.

Results

Tables, which you have permissions to, appear for the specified schema. The list includes both local and virtual tables. Virtual tables are those with the green, remote sources icon ()

13.1.12.1.4 Delete a Virtual Table

Remove an existing virtual table from the local system using SAP HANA studio.

Prerequisites

One of the following:

- You created the virtual table.
- You have the DROP privilege on the object created by another user.

Context

If access to a virtual table has been granted to multiple users, deleting the table removes it for all users, not just the user deleting it.

Procedure

1. In the *Systems* view, expand the folders to display the table to be dropped: ► *Catalog* ► *<schema_name>* ► *Tables* ▼.
2. Right-click the virtual table and choose *Delete* from the context menu.
If there are dependent virtual tables, a warning appears listing the dependent tables and warning that they will also be deleted.
3. Choose *OK* to confirm.

13.1.12.2 Managing Virtual Tables Using SQL Console

Create, view, refresh, and delete virtual tables using SQL.

13.1.12.2.1 Create a Virtual Table

Create a virtual table using a SQL console.

Prerequisites

One of the following:

- You created the remote source.
- You have the CREATE VIRTUAL TABLE object privilege on the remote source created by another user.

Procedure

1. Open a SQL console.
2. Execute the command specifying the schema (owner) and name of the virtual table and the name of the :

```
CREATE VIRTUAL TABLE [<schema>.]<virtual_table_name> AT  
    "<remote_source>".<database_name>.<schema>.<table_name>;
```

❁ Example

The following statement creates a virtual table of table T1 in the remote source HANA1. The remote database name is DT1 and the owner of T1 is user2:

```
CREATE VIRTUAL TABLE HANA1_T1 AT HANA1.DT1.user2.T1;
```

13.1.12.2.2 List Virtual Tables by Schema

Display the virtual tables of a remote source by schema using an SQL console.

Context

You cannot change the definition or contents while viewing.

Procedure

1. Open a SQL console.
2. Execute the `SELECT` statement specifying the schema name of the virtual tables to list.

```
SELECT * FROM "SYS"."VIRTUAL_TABLES" WHERE SCHEMA_NAME='<schema_name>'
```

Tables, which you have permissions to, appear for the specified schema. The list includes both local and virtual tables.

Example

The following statement lists the tables for schema `user1`:

```
SELECT * FROM SYS.VIRTUAL_TABLES WHERE SCHEMA_NAME='<user1>'
```

13.1.12.2.3 Refresh Virtual Tables

Update a virtual table to reflect metadata changes in the corresponding remote source table using an SQL console.

Prerequisites

- Requires the `CREATE VIRTUAL TABLE` system privilege.

Context

When changes to the metadata in a remote table are made, the changes are not automatically reflected in the corresponding virtual table. Manually update the virtual table to reflect the changes.

Procedure

To refresh the virtual table, execute:

```
ALTER VIRTUAL TABLE [<schema_name>.] "<virtual_table_name>" REFRESH DEFINITION
```

❖ Example

The following statement refreshes the content of virtual table HANA1_T1 in schema user1:

```
ALTER VIRTUAL TABLE user1."HANA1_T1" REFRESH DEFINITION
```

13.1.12.2.4 Delete a Virtual Table

Remove an existing virtual table from the target system using a SQL console.

Prerequisites

One of the following:

- You created the virtual table.
- You have the DROP privilege on the object created by another user.

Context

If access to a virtual table has been granted to multiple users, deleting the table removes it for all users, not just the user deleting it.

Procedure

1. Open a SQL console.
2. Execute:

```
DROP TABLE [<schema_name>]. "<virtual_table_name>";
```

❖ Example

The following statement deletes virtual table HANA1_T1 in the schema user1:

```
DROP TABLE user1."HANA1_T1" CASCADE;
```

13.1.12.3 Managing Virtual Tables Using SAP HANA Database Explorer

Create, view table content and definition, and delete virtual tables using SAP HANA Database Explorer

13.1.12.3.1 Create a Virtual Table

Create a virtual table from the remote object of a remote source.

Prerequisites

One of the following:

- You created the remote source.
- You have the CREATE VIRTUAL TABLE object privilege on the remote source created by another user.

Procedure

1. Right-click your remote source in the catalog browser item list and click *Open*.
The remote source editor opens.
2. Click on the *Remote Objects* tab to view all remote objects associated with your remote source.
3. Click one or more tables and click  to open the *Create Virtual Objects* dialog.
4. Specify a schema to add the virtual object to, then specify an object name prefix for the virtual objects you are creating.
5. Click *Create*.
6. (Optional) Repeat steps 1-5 for each schema that you want to add your virtual table to.

Results

Your virtual table is created. When you open the catalog browser for the schema you added it to and click on *Tables*, you can see the virtual table in the catalog browser list.

13.1.12.3.2 Delete a Virtual Table

Delete an existing virtual table from your schema using SAP HANA Database Explorer.

Prerequisites

One of the following:

- You created the virtual table.
- You have the DROP privilege on the object created by another user.

Context

If access to a virtual table has been granted to multiple users, deleting the table removes it for all users, not just the user deleting it.

Procedure

Right-click the virtual table in the catalog browser item list and click *Delete*.

Results

The virtual table is deleted from your schema.

13.1.12.3.3 View the Definition and Contents of a Virtual Table

Display the definition and contents of a virtual table using SAP HANA Database Explorer.

Context

You cannot change the definition or the contents of the table while viewing.

Procedure

1. In the catalog browser, click [Tables](#).
2. In the catalog browser item list, right-click your virtual table (virtual tables are those with the green, remote sources icon) and click [Open](#)

13.1.13 Managing Linked Database

Management of linked database can only be performed using SQL syntax.

i Note

Managing linked database is not applicable when optimized mode is enabled.

13.1.13.1 Drop Linked Tables

Drop all internally generated objects associated with linked database using the remote source.

Prerequisites

- Requires the ALTER object privilege on the remote source.

Context

The CASCADE option drops all the linked tables and dependent objects associated with linked tables. The RESTRICT option drops linked tables only if there are no dependencies on any of the linked tables. If this option is used and there are dependent objects on a linked table, an error is raised, and all linked tables are retained. If no drop option is specified, all internally generated linked tables with no dependencies are dropped. Linked tables with dependencies are retained.

Procedure

To clear generated objects, execute:

```
ALTER REMOTE SOURCE <remote_source_name> DROP LINKED OBJECTS [CASCADE | RESTRICT]
```

❖ Example

To clear linked objects associated with the remote source `myremotesys` that have no current references to a linked object, execute:

⇐ Sample Code

```
ALTER REMOTE SOURCE myremotesys DROP LINKED OBJECTS
```

To clear linked objects associated with the remote source `myremotesys` and drop any dependent objects that reference the linked object, execute:

⇐ Sample Code

```
ALTER REMOTE SOURCE myremotesys DROP LINKED OBJECTS CASCADE
```

13.1.13.2 Refresh Linked Database

Refresh metadata for a single linked table or all linked objects using the remote source.

Prerequisites

- Requires the ALTER object privilege on the remote source.

Context

The REFRESH LINKED OBJECTS clause refreshes metadata for all linked objects. Definitions of any tables changed on the remote source are updated. Refreshing metadata could potentially take some time to execute. It is recommended that you schedule the execution within an appropriate maintenance window. The REFRESH LINKED TABLE clause refreshes metadata for a single specified table. Use this clause to perform an immediate refresh of a changed table when it would be difficult to execute the refresh on the entire linked database.

When refreshing a linked table, `<table_name>` is the name of the table on the remote source.

Procedure

To refresh metadata on all linked objects, execute:

```
ALTER REMOTE SOURCE <remote_source> {REFRESH LINKED OBJECTS  
| REFRESH LINKED TABLE <remote_source>.<schema_name>.<table_name>}
```

❖ Example

To refresh metadata of all linked objects associated with remote source `myremotesys`, execute:

≡ Sample Code

```
ALTER REMOTE SOURCE myremotesys REFRESH LINKED OBJECTS
```

To refresh metadata for remote table `myschema.mytable` on remote source `myremotesys`, execute:

≡ Sample Code

```
ALTER REMOTE SOURCE myremotesys REFRESH LINKED TABLE  
myremotesys.myschema.mytable
```

13.1.14 EXPORT/IMPORT Virtual Tables

Use EXPORT/IMPORT statements to move data between systems and for troubleshooting purposes.

i Note

- Exporting and importing virtual tables is not applicable when using linked database.

13.1.14.1 EXPORT/IMPORT Virtual Tables Between Systems

Export virtual tables from one system and import them into another.

Prerequisites

- A remote source must exist on the target system using the same name as the source system.
- You have the system privileges IMPORT, EXPORT, and INSERT.

Context

You can export some or all of the existing virtual tables in the source system to a new system, specifying a new schema or database name during the process.

Both the source and target systems must be running SAP HANA 2.0 SPS 01 or later, but do not have to be running the same version.

Ensure you save the export file to a location available to the target system.

EXPORT/IMPORT When Schema Name Is the Same on Target and Source Systems

Procedure

1. On the source system, export part or all of the schema to a temporary directory.

```
EXPORT <schema_name>.[ <table_name> | "*" ] AS CSV INTO '<temporary_dir>';
```

2. Create the remote source on the target system using the same name as the source system.
3. On the target system, import the schema.

```
IMPORT "<schema_name>".[<table_name> | * ] AS CSV FROM '<temporary_dir>';
```

❖ Example

Export all tables in `MYSCHEMA` to remote source `RS`, using a temporary directory `/export_hold`.

```
EXPORT MYSCHEMA.* AS CSV INTO '/export_hold';
CREATE REMOTE SOURCE RS ADAPTER "hanaodbc" CONFIGURATION
'Driver=libodbcHDB.so;ServerNode=mymachine:30115'
WITH CREDENTIAL TYPE 'PASSWORD' USING
'user=<user_name>;password=<password>';
IMPORT "MYSCHEMA".*" AS CSV FROM '/export_hold';
```

Export only table `VT` in `MYSCHEMA` to remote source `RS`, using a temporary directory `/export_hold`.

```
EXPORT MYSCHEMA.VT AS CSV INTO '/export_hold';
CREATE REMOTE SOURCE RS ADAPTER "hanaodbc" CONFIGURATION
'Driver=libodbcHDB.so;ServerNode=mymachine:30115'
WITH CREDENTIAL TYPE 'PASSWORD' USING
'user=<user_name>;password=<password>';
IMPORT "MYSCHEMA"."VT" AS CSV FROM '/export_hold';
```

Rename the Remote Object During Transport

Procedure

1. On the source system, export part or all of the schema to a temporary directory.

```
EXPORT <schema_name>.[ <table_name> | "*" ] AS CSV INTO '<temporary_dir>';
```

2. Create the remote source on the target system using the same name as the source system.

3. On the target system, import the schema, specifying the `<original_database_name>` and `<original_schema_name>` and the `<new_database_name>` and `<new_schema_name>`.

```
IMPORT "<schema_name>". "[<table_name> | * ]" AS CSV FROM '<temporary_dir>'
WITH RENAME REMOTE OBJECT
'<remote_source>.<original_database_name>.<original_source_schema>.TBL' TO
'<remote_source>.<new_database_name>.<new_source_schema>.TBL';
```

❁ Example

Export table VT in MYSCHEMA to remote source RS, with a database name ADMIN and a schema name RS_SCHEMA.

```
EXPORT MYSCHEMA."VT" AS CSV INTO '/export_hold';
CREATE REMOTE SOURCE RS ADAPTER "hanaodbc" CONFIGURATION
'Driver=libodbcHDB.so;ServerNode=mymachine:30115'
WITH CREDENTIAL TYPE 'PASSWORD' USING
'user=<user_name>;password=<password>';
IMPORT "MYSCHEMA"."VT" AS CSV FROM '/tmp' WITH RENAME REMOTE OBJECT
'RS.ADMIN.RS_SCHEMA.TBL' TO 'RS.NEW_ADMIN.NEW_RS_SCHEMA.TBL';
```

13.1.14.2 EXPORT/IMPORT Virtual Tables for Debugging Purposes

Reproduce a virtual table workflow in a local environment for troubleshooting purposes.

Prerequisites

You have the system privileges IMPORT, EXPORT, and INSERT.

Context

During import, a loopback remote source is created that points to the local SAP HANA server. For every virtual table, a corresponding dummy local table is created with the same metadata. Virtual tables are then created using this loopback remote source and the corresponding dummy local table.

Procedure

1. On the source system, export part or all of the schema to a temporary directory.

```
EXPORT <schema_name>.[ <table_name> | "*" ] AS CSV INTO '<temporary_dir>';
```

2. On the target system, import the schema using a loopback remote source.

```
IMPORT "<schema_name>"."<table_name> | * )" AS CSV FROM '<temporary_dir>'
WITH LOOPBACK REMOTE SOURCE;
```

❖ Example

Export all tables in MYSCHEMA to remote source RS, using a temporary directory /export_hold.

```
EXPORT MYSCHEMA."*" AS CSV INTO '/export_hold';
IMPORT "MYSCHEMA"."*" AS CSV FROM '/export_hold' WITH LOOPBACK REMOTE SOURCE;
```

Export only table VT in MYSCHEMA to remote source RS, using a temporary directory /export_hold.

```
EXPORT MYSCHEMA.VT AS CSV INTO '/export_hold';
IMPORT "MYSCHEMA"."VT" AS CSV FROM '/export_hold' WITH LOOPBACK REMOTE SOURCE;
```

13.1.15 Monitor Remote Connections and Statements

Monitor active connections and running statements on remote connections using SAP HANA studio or SAP HANA cockpit.

Prerequisites

- Requires the CATALOG READ privilege to view the details.

Context

Use the monitoring tools to monitor:

Remote connections active in the database

This tool provides details about the connections that were opened in the current session, including when the connection was opened, how many remote statements were executed, and the name of the remote source.

Remote statements executed in the database

This tool allows you to see the full SQL text of the SQL statements executed on remote sources. It also shows you when the query was started, how long the query took, and the number of records that were returned.

Monitor Remote Connections using SAP HANA Cockpit

Context

The number of running statements and active connections appears under Smart Data Access. Click the information type to display details,

Procedure

1. On the system overview, scroll to *SAP HANA Smart Data Access Administration*.
2. Choose the information type to monitor.
 - Running Statements
 - Active Connections

Monitor Remote Connections using HANA Studio

Procedure

1. In the *Systems* view, expand your system's *Provisioning* node.
2. Select *Smart Data Access* and from the context menu choose *Open Smart Data Access Administration*.
3. Choose the appropriate tab:
 - Query Monitoring
 - Connection Monitoring

13.1.15.1 Monitor Details Using SAP HANA Studio

Detailed information about the remote statements executed and active connections active in the database.

Context

Use the monitoring tools to monitor:

Remote connections active in the database

Provides details about the connections that were opened in the current session, including when the connection was opened, how many remote statements were executed, and the name of the remote source.

Remote statements executed in the database Allows you to see the full SQL text of the SQL statements executed on remote sources. It also shows you when the query was started, how long the query took, and the number of records that were returned.

i Note

This functionality is not currently available using SAP HANA Database Explorer.

Procedure

1. In the *Systems* view, expand your system's *Provisioning* node.
2. Select *Smart Data Access* and from the context menu choose *Open Smart Data Access Administration*.
3. Choose the appropriate tab:
 - Query Monitoring
 - Connection Monitoring

Results

Information available for remote connections.

Detail	Description
Connection	Connection ID
Status	Connection status: <ul style="list-style-type: none"> • <i>Connected</i>: Connection is active • <i>Disconnected</i>: Connection has been closed
Client	Client host name
Source Name	Name of the remote data source
Source User	Name of the remote data source user
Start Time	Start time of first query execution
Statements	Number of statements executed
Details	Connection details, including, for example, the data source name and DML_MODE

Information available for remote statements

Detail	Description
SQL Statement	Full SQL string
Start Time	Start time of query execution
End Time	End time of query execution
SAP HANA Studio: Execution Time (ms)	Query execution time

Detail	Description
SAP HANA cockpit: Statement Runtime (Seconds)	
Status	Query execution status: <ul style="list-style-type: none"> Analyzing: Query is being analyzed by the query optimizer Optimizing: Query is being optimized by the query optimizer Executing: Query is running Closed: Query has completed Failed: Query execution failed
Rows	Number of rows returned in the query result
Fetch Size	Specifies the byte size of fetched records
Remote Source Name	Name of the remote data source

13.1.15.2 Monitor Details Using SAP HANA Cockpit

Detailed information about the remote statements executed and active connections active in the database.

Context

Use the monitoring tools to monitor:

- Remote connections active in the database** Provides details about the connections that were opened in the current session, including when the connection was opened, how many remote statements were executed, and the name of the remote source.
- Remote statements executed in the database** Allows you to see the full SQL text of the SQL statements executed on remote sources. It also shows you when the query was started, how long the query took, and the number of records that were returned.

i Note

This functionality is not currently available using SAP HANA Database Explorer.

Procedure

- On the system overview, scroll to *SAP HANA Smart Data Access Administration*.
- Choose the information type to monitor.
 - Running Statements

- Active Connections

Results

Information available for remote connections

Detail	Description
Connection	Connection ID
Adapter	Name of the adapter used for Smart Data Access
Status	Connection status: <ul style="list-style-type: none"> • <i>Connected</i>: Connection is active • <i>Disconnected</i>: Connection has been closed
Source Name	Name of the remote data source
Source User	Name of the remote data source user
Start Time	Start time of first query execution
Statements	Number of statements executed
Details	Connection details, including, for example, the data source name and DML_MODE

Information available for remote statements

Detail	Description
SQL Statement	Full SQL string
Start Time	Start time of query execution
End Time	End time of query execution
SAP HANA Studio: Execution Time (ms)	Query execution time
SAP HANA cockpit: Statement Runtime (Seconds)	
Status	Query execution status: <ul style="list-style-type: none"> • Analyzing: Query is being analyzed by the query optimizer • Optimizing: Query is being optimized by the query optimizer • Executing: Query is running • Closed: Query has completed • Failed: Query execution failed
Rows	Number of rows returned in the query result
Fetch Size	Specifies the byte size of fetched records
Remote Source Name	Name of the remote data source
User	User who executed the statement
Transaction	Transaction ID

13.1.15.3 Monitor Details Using SQL Console

Detailed information about the remote statements executed and active connections active in the database.

Context

Use the monitoring tools to monitor:

Remote connections active in the database Provides details about the connections that were opened in the current session, including when the connection was opened, how many remote statements were executed, and the name of the remote source.

Remote statements executed in the database Allows you to see the full SQL text of the SQL statements executed on remote sources. It also shows you when the query was started, how long the query took, and the number of records that were returned.

i Note

This functionality is not currently available using SAP HANA Database Explorer.

Procedure

1. Open a SQL console.
2. For information on remote statements, execute a SELECT statement using the M_REMOTE_STATEMENTS system view.

```
SELECT * FROM SYS.M_REMOTE_STATEMENTS
```

3. For information on connections, execute a SELECT statement using the M_REMOTE_CONNECTIONS system view.

```
SELECT * FROM SYS.M_REMOTE_CONNECTIONS
```

Results

Details	Description
CONNECTION_ID	Specifies the connection ID
TRANSACTION_ID	Specifies the transaction ID
STATEMENT_ID	Specifies the HANA statement ID

Details	Description
REMOTE_CONNECTION_ID	Specifies the ID of the remote connection
REMOTE_SOURCE_NAME	Specifies the remote source name
START_TIME	Specifies the statement start time
END_TIME	Specifies the statement end time
FETCHED_RECORD_COUNT	Specifies the number of fetched records
FETCHED_SIZE	Specifies the byte size of fetched records
REMOTE_STATEMENT_STATUS	Specifies the statement status: EXECUTING, CLOSED
REMOTE_STATEMENT_STRING	Specifies the statement string
USER_NAME	Specifies the user name
REMOTE_STATEMENT_DETAILS	Specifies the statement details

Details	Description
CONNECTION_ID	Specifies the connection ID
TRANSACTION_ID	Specifies the transaction ID
STATEMENT_ID	Specifies the HANA statement ID
REMOTE_CONNECTION_ID	Specifies the ID of the remote connection
REMOTE_SOURCE_NAME	Specifies the remote source name
START_TIME	Specifies the statement start time
END_TIME	Specifies the statement end time
FETCHED_RECORD_COUNT	Specifies the number of fetched records
FETCHED_SIZE	Specifies the byte size of fetched records
REMOTE_STATEMENT_STATUS	Specifies the statement status: EXECUTING, CLOSED
REMOTE_STATEMENT_STRING	Specifies the statement string
USER_NAME	Specifies the user name
REMOTE_STATEMENT_DETAILS	Specifies the statement details

13.1.16 Data Type Support

Data type support varies by remote source.

The SAP HANA Smart Data Access 2.0 master note lists by remote source how data types are mapped to SAP HANA. Data types not listed are not supported.

Related Information

[SAP Note 2352696](#)

13.1.16.1 Spatial Data Types

Smart data access supports virtual tables on an SAP HANA remote source containing spatial data types.

The ST_POINT data type in the remote source maps to the ST_GEOMETRY data type in the virtual table.

Spatial data type support has some functional restrictions. See [2609914 - Smart Data Access - Functional Restrictions for Supporting Spatial Data Types](#) for a list of these functional restrictions.

For information on SAP HANA spatial data types, see *SAP HANA Spatial Reference*

13.1.17 Functions Pushed Down to Remote Sources

When executing queries on virtual tables that use SAP HANA functions, whenever possible smart data access pushes execution of the function to the remote source to improve query performance.

The ability to push an SAP HANA function down to a remote source depends on the capabilities of the remote source. In some instances, where the SAP HANA function does not map one-to-one to a remote source function, but an equivalent remote function is available, the equivalent remote function is used. If a remote source doesn't support an SAP HANA function, the SAP HANA function is executed on the local host.

The SAP HANA Smart Data Access 2.0 master note lists by remote source the SAP HANA functions that can be pushed down, and where applicable, the equivalent remote source function. Where the SAP HANA function does not map one-to-one, the list indicates what equivalent function is pushed down.

Related Information

[SAP Note 2352696 - SAP HANA Smart Data Access 2.0 Master Release Note](#)

13.1.18 Synonyms

Use synonyms to create an alternative name for a virtual table or a remote table when using linked database.

When creating the synonym for linked database, specify the three part name. For virtual tables, specify just `<identifier>` or `<schema_name>.<identifier>`.

```
CREATE SYNONYM <synonym_name> FOR [[<database_name>.<schema_name>.<identifier>
```

where:

Variable	Description
<code><database_name></code>	Specifies the name of the remote source for linked database. This value is not applicable to virtual tables.
<code><schema_name></code>	For linked database, specifies the schema name of the table on the remote source. For virtual tables, specifies the local schema and virtual table name.
<code><identifier></code>	For linked database, specifies the name of the table on the remote source. For virtual tables, specifies the name of the local virtual table.

Example

Create synonym `table1_synonym` for table `admin.table1` on remote source `source1`.

Sample Code

```
CREATE table1_synonym FOR source1.admin.table1
```

Create synonym `table2_synonym` for virtual table `admin.table2` on remote source `source2`.

Sample Code

```
CREATE table2_synonym FOR admin.table2
```

13.1.19 Statistics on Virtual Tables and Linked Database

Statistics assist the query optimizer in making better decisions and work for both virtual tables and linked database.

13.1.19.1 Create Statistics on a Virtual Table or Linked Database

Create data statistic virtual objects that the query optimizer uses to make better decisions for query plans.

Prerequisites

One of the following:

- You created the virtual table you are creating statistics on.
- You have the ALTER privilege on the object you are creating statistics on.
- For linked database, you require the LINKED DATABASE object level privilege on the remote source, regardless of who created the remote source.

Context

See *Monitor Remove Connections and Statement* to check that the statement executed correctly.

Procedure

To create statistics on a virtual table or linked database, execute:

```
CREATE STATISTICS <data_statistics_name> ON <data_sources>
  <data_statistics_type>
  [ <data_statistics_properties> ]
  [ <initial_refresh> ]
```

<data_statistics_name>

Specifies a unique name for the data statistics object.

```
<data_statistics_name> ::= [<schema_name>.]<identifier>
```

<data_statistics_name> is only allowed when the result of the creation is a single data statistics object. The number of data statistics objects created by CREATE STATISTICS is determined by the combination of <data_statistics_type> and the

number of columns specified in `<data_sources>.<data_statistics_name>` is not supported for linked database.

`<data_sources>`

Specifies the remote data source you want to create data statistics objects for.

```
<data_sources> ::= { <virtual_table_name> |  
<linked_database_table_name> }  
<virtual_table_name> ::=  
[<schema_name>.]<virtual_table_name> [(<column_name>{,  
<column_name>}...)] ]  
<linked_database_table_name> ::=  
<remote_source_name>.<schema_name>.<remote_table_name>  
[(<column_name>{, <column_name>}...)] ]
```

For RECORD COUNT statistics objects, you cannot specify columns as part of `<virtual_table_name>` or `<linked_database_table_name>`.

- `<schema_name>` Specifies the name of the schema in the remote source.
- `<virtual_table_name>` Specifies the name of the remote table.
- `<column_name>` Specifies a table column in the remote source.
- `<remote_source_name>` Specifies the name of the remote source.
- `<remote_table_name>` Specifies the name of the table in the remote source

`<data_statistics_type>`

Specifies the data statistics type.

```
<data_statistics_type> := TYPE <type_name>  
<type_name> ::=  
HISTOGRAM  
| SIMPLE  
| TOPK  
| SKETCH  
| SAMPLE  
| RECORD COUNT
```

A data source can have only one data statistics object of a certain type. For example, column A of table T can have one data statistics object of type HISTOGRAM and one of type SIMPLE. If the TYPE clause is not specified, then the default is HISTOGRAM. Some data statistic types may not be appropriate for a given data source.

HISTOGRAM Creates a data statistics object that helps the query optimizer estimate the data distribution in a single-column data source. If you specify multiple columns in `<data_sources>`, then multiple data statistics objects (histograms) are created--one per column specified.

SIMPLE Creates a data statistics object that helps the query optimizer calculate basic statistics, such as min, max, null count, count, and distinct count for a single-column data source. If you specify multiple columns in `<data_sources>`, then multiple data statistics objects are created--one per column specified. When beneficial, the SQL

optimizer maintains system SIMPLE data statistics objects automatically on column and row store tables only.

TOPK Creates a data statistics object that helps the query optimizer identify the highest-frequency values in a table data source. If you specify multiple columns in `<data_sources>`, then multiple data statistics objects are created--one per column specified. When beneficial, the SQL optimizer maintains system TOPK data statistics objects automatically on column and row store tables only.

SKETCH Creates a data statistics object that helps the query optimizer estimate the number of distinct values in the data source. A data statistics object is created for the specified `<table_name> (<column-name>, ...)`, which approximates the number of distinct tuples in the projection of the table on the set of specified columns.

SAMPLE Creates samples of data from `<data_source>` that the SQL optimizer can use during optimization. When beneficial, the SQL optimizer generates system SAMPLE data statistics objects automatically on column store tables only. However, this behavior can incur a cost to performance. You can avoid this cost by creating SAMPLE data statistics objects explicitly (in advance). Creating them explicitly is especially useful in situations where sampling live table data is expensive (for example, very big tables).

RECORD COUNT Creates a data statistics object that helps the query optimizer calculate the number of records (rows) in a table data source. The RECORD COUNT type is a table-wide statistic that can be created on all table types except column store tables. You do not specify columns in `<data_sources>` when creating a record count data statistics object. When beneficial, the SQL optimizer maintains system RECORD COUNT data statistics objects automatically on column store tables only.

`<data_statistics_properties>`

Specifies the properties of the data statistics object.

```
<data_statistics_properties> ::=
  <data_statistics_property> [<data_statistics_property>]...
<data_statistics_property> ::=
  | REFRESH TYPE <refresh_type>
  | ENABLE <on_off>
  | BUCKETS <unsigned_integer>
  | QERROR <numeric_literal>
  | QTHETA <unsigned_integer>
  | { MEMORY <memory_bytes> | MEMORY PERCENT
    <memory_percentage> }
  | ACCURACY <numeric_literal>
  | PREFIX BITS <unsigned_integer>
  | PERSISTENT <on_off>
  | VALID FOR <valid_for_list>
  | CONSTRAINT '<constraint_param>'
```

Restrictions to which properties apply to which statistic types are noted in the property descriptions.

REFRESH TYPE

`<refresh_type>`

Specifies the strategy for refreshing the data statistics object.

```
<refresh_type> ::= MANUAL | DEFAULT
```

MANUAL specifies that the database statistics object is not refreshed until a rebuild is explicitly requested by a REFRESH STATISTICS statement.

DEFAULT specifies that the database server decides the best refresh strategy based on the data source.

REFRESH TYPE only affects data statistics objects that are enabled.

ENABLE `<on_off>`

Controls whether the optimizer uses the data statistics object.

```
<on_off> ::= ON | OFF
```

ENABLE ON enables the optimizer to use of the data statistics object. ENABLE ON specified with NO INITIAL REFRESH returns an error. However, the data statistics object must also be populated with data (refreshed) for the optimizer to use it. The default behavior is ENABLE ON.

ENABLE OFF disables the use of the data statistics object by the optimizer and prevents the ability to refresh the data statistics object. Data statistics objects that are not enabled can still be dropped. To make a data statistics object with ENABLE OFF accessible to the optimizer, execute an ALTER STATISTICS...ENABLE ON statement.

BUCKETS

`<unsigned_integer>`

It is only for use with TYPE HISTOGRAM or TOPK. For HISTOGRAM, BUCKETS specifies the maximum number of data buckets in the histogram. For TOPK, BUCKETS specifies the K value.

The default is automatically determined by the data statistics building algorithm in use.

ACCURACY

`<numeric_literal>`

Controls the time and space requirements to use for the sketch algorithms. This parameter can only be specified when TYPE is SKETCH and must be a number between 0 and 1, with larger values causing decreased time and space requirements but poorer sketch resolution. The default is 0.1.

PREFIX BITS

`<unsigned_integer>`

Controls the number of bits the sketch algorithms use when constructing the sketch statistics. Specify this

parameter when TYPE is SKETCH, and its value is an integer between 0 and 63. The default is 8.

CONSTRAINT

`<constraint_param>`

Specifies constraints to use for the specified `<data_statistics_type>`

- For HISTOGRAM, `<constraint_param>` specifies the mathematical constraint for the histogram:

```
<constraint_param> ::= MAXDIFF
```

A non-default CONSTRAINT for histograms results in an error.

- For SKETCH, `<constraint_param>` specifies the algorithm to use to build the sketch. The default is LOGLOGCOUNTING; the remaining algorithms are for internal use.

```
<constraint_param> ::=  
KMINVAL  
| PCSA  
| LINEARCOUNTING  
| LOGCOUNTING  
| LOGLOGCOUNTING  
| SUPERLOGLOGCOUNTING
```

`<initial_refresh>`

Specifies whether to populate the data statistics object with data after creation.

```
<initial_refresh> ::= [ NO ] INITIAL REFRESH
```

**INITIAL
REFRESH**

Creates the definition of the data statistics object and populates it with data. The default behavior is INITIAL REFRESH.

**NO INITIAL
REFRESH**

Creates the definition of the data statistics object, but does not populate it with data.

Use NO INITIAL REFRESH when you want to change the underlying data before refreshing the data statistics object.

Examples

Create SIMPLE statistics on virtual table HANA1_T1.

```
CREATE STATISTICS "TEST1" ON HANA1_T1 (A1) TYPE SIMPLE;
```

Using linked database, create TOPK statistics on table T1 using remote source HANA1 and schema MYSCHEMA with 10 buckets.

```
CREATE STATISTICS ON HANA1.MYSCHEMA.T1 (A1) TYPE TOPK BUCKETS 10;
```

Related Information

[Monitor Details Using SAP HANA Cockpit \[page 1937\]](#)

[Monitor Details Using SAP HANA Studio \[page 1935\]](#)

[Monitor Details Using SQL Console \[page 1939\]](#)

[Alter Statistics on a Virtual Table or Linked Database \[page 1948\]](#)

13.1.19.2 Alter Statistics on a Virtual Table or Linked Database

Alter the properties of a data statistic object for virtual tables or linked database.

Prerequisites

One of the following:

- You created the virtual table you are altering statistics on.
- You have the ALTER privilege on the object you are altering statistics on.
- For linked database, you require the LINKED DATABASE object level privilege on the remote source, regardless of who created the remote source.

Procedure

To alter statistics on a virtual table or linked database, execute:

```
ALTER STATISTICS <data_statistics_name> ON <data_sources>  
<data_statistics_type>[[HAVING] <match_properties>]]  
[ <set_data_statistics_properties>]  
[<initial_refresh>]
```

<data_statistics_name>

Specifies the name of the data statistics object to alter.

```
<data_statistics_name> ::= [<schema_name>.]<identifier>
```

<data_sources>

Specifies the remote data source(s) of the data statistics objects to alter.

```
<data_sources> ::= { <virtual_table_name> |  
<linked_database_table_name> }  
<virtual_table_name> ::=  
[<schema_name>.]<virtual_table_name> [(<column_name>{,  
<column_name>}...)] ] ]  
<linked_database_table_name> ::=
```

```
<remote_source_name>.<schema_name>.<remote_table_name>  
[( <column_name>{, <column_name>}... ) ] ]
```

- <schema_name>** Specifies the name of the schema in the remote source.
- <virtual_table_name>** Specifies the name of the remote table.
- <column_name>** Specifies a table column in the remote source.
- <remote_source_name>** Specifies the name of the remote source.
- <remote_table_name>** Specifies the name of the table in the remote source

<match_properties>

Specifies properties to use for matching when selecting data statistics to alter.

```
<match_properties> ::= <match_property> [<...>]  
<match_property> ::=  
  TYPE <data_statistics_type>  
  | REFRESH TYPE <refresh_type_filter>
```

If TYPE is not specified, then all data statistics objects on specified data sources are altered (ALL). See the CREATE STATISTICS Statement topic for descriptions of the supported data statistics types.

- <data_statistics_type>** Specifies the type of data statistics objects to match when selecting the data statistics to alter.

```
<data_statistics_type> := TYPE  
<type_name>  
<type_name> ::=  
  HISTOGRAM  
  | SIMPLE  
  | TOPK  
  | SKETCH  
  | SAMPLE  
  | RECORD COUNT  
  | ALL
```

See the Create Statistics on a Virtual Table or Linked Database topic for descriptions of the supported data statistics types.

- <refresh_type_filter>** Specifies the refresh strategy to match in the data statistics objects when selecting the data statistics to alter. ALL is the default.

```
<refresh_type_filter> ::= MANUAL | ALL
```

<set_data_statistics_properties>

Specifies the properties of the data statistics objects you set using the SET keyword.

```
SET <data_statistics_properties>  
<data_statistics_properties> ::=  
<data_statistics_property> [<data_statistics_property>]...  
  
<data_statistics_property> ::=  
  | REFRESH TYPE <refresh_type>
```

```
| ENABLE <on_off>
| BUCKETS <unsigned_integer>
| ACCURACY <numeric_literal>
| PREFIXBITS <unsigned_integer>
| PERSISTENT <on_off>
| CONSTRAINT <constraint_param>
```

REFRESH TYPE
<refresh_type>

Specifies the strategy for refreshing the data statistics object.

```
<refresh_type> ::= MANUAL | DEFAULT
```

MANUAL specifies that the database statistics object is not refreshed until a rebuild is explicitly requested by a REFRESH STATISTICS statement.

DEFAULT specifies that the database server decides the best refresh strategy based on the data source.

REFRESH TYPE only affects data statistics objects that are enabled.

ENABLE <on_off>

Controls whether the optimizer uses the data statistics object.

```
<on_off> ::= ON | OFF
```

ENABLE ON enables the optimizer to use of the data statistics object. ENABLE ON specified with NO INITIAL REFRESH returns an error. However, the data statistics object must also be populated with data (refreshed) for the optimizer to use it. The default behavior is ENABLE ON.

ENABLE OFF disables the use of the data statistics object by the optimizer and prevents the ability to refresh the data statistics object. Data statistics objects that are not enabled can still be dropped. To make a data statistics object with ENABLE OFF accessible to the optimizer, execute an ALTER STATISTICS...ENABLE ON statement.

BUCKETS
<unsigned_integer>

It is only for use with TYPE HISTOGRAM or TOPK. For HISTOGRAM, BUCKETS specifies the maximum number of data buckets in the histogram. For TOPK, BUCKETS specifies the K value.

The default is automatically determined by the data statistics building algorithm in use.

ACCURACY
<numeric_literal>

Controls the time and space requirements to use for the sketch algorithms. This parameter can only be specified when TYPE is SKETCH and must be a number between 0 and 1, with larger values causing decreased time and space requirements but poorer sketch resolution. The default is 0.1.

PREFIX BITS**<unsigned_integer>**

Controls the number of bits the sketch algorithms use when constructing the sketch statistics. Specify this parameter when TYPE is SKETCH, and its value is an integer between 0 and 63. The default is 8.

CONSTRAINT**<constraint_param>**

Specifies constraints to use for the specified <data_statistics_type>. For SKETCH, <constraint_param> specifies the algorithm to use to build the sketch. The default is LOGLOGCOUNTING; the remaining algorithms are for internal use.

```
<constraint_param> ::=
KMINVAL
| PCSA
| LINEARCOUNTING
| LOGCOUNTING
| LOGLOGCOUNTING
| SUPERLOGLOGCOUNTING
```

<initial_refresh>

Specifies whether to repopulate the data statistics object with data after altering it.

```
<initial_refresh> ::= [ NO ] INITIAL REFRESH
```

If the object was built, then disabled, and is now being re-enabled, initial refresh is not required

INITIAL REFRESH

Alters the definition of the data statistics object and repopulates it with data. The default behavior is INITIAL REFRESH.

NO INITIAL REFRESH

Alters the definition of the data statistics object, but does not repopulate it with data.

Use NO INITIAL REFRESH when you want to change the underlying data before refreshing the data statistics object.

Examples

The following example sets the number of buckets to 150 for the virtual table Remote1_A1.

```
ALTER STATISTICS on MYSYSTEM.REMOTE2_A1 TYPE TOPK SET BUCKETS 10 NO INITIAL REFRESH;
```

The following example sets the number of buckets to 10 on the remote source remote2 using linked database.

```
ALTER STATISTICS on "remote2"."SYSTEM"."A1" TYPE TOPK SET BUCKETS 10;
```

Related Information

[Monitor Details Using SAP HANA Cockpit \[page 1937\]](#)

[Monitor Details Using SAP HANA Studio \[page 1935\]](#)

13.1.19.3 Refresh Statistics on a Virtual Table or Linked Database

Refreshes data statistic virtual objects that the query optimizer uses to make better decisions for query plans.

Prerequisites

One of the following:

- You created the virtual table you are refreshing statistics on.
- You have the ALTER privilege on the object you are refreshing statistics on.
- For linked database, you require the LINKED DATABASE object level privilege on the remote source, regardless of who created the remote source.

Context

Specify EXACT to refresh a data statistics virtual object that precisely matches `<data_sources>` (including column order). Specify CASCADE to refresh data statistics objects that reference at least one column in `<data_sources>`. If `<match_type>` is not specified, then any data statistics objects that reference all or some of the columns specified in `<data_sources>` are refreshed.

Procedure

1. To refresh statistics on a virtual table, execute:

```
REFRESH STATISTICS ON [<schema_name>.]<virtual_table_name>  
[ [HAVING] <match_properties>
```

2. Using linked database, to refresh statistics on a remote table, execute:

```
REFRESH STATISTICS <database_name>.<schema_name>.<remote_table_name>  
[ [HAVING] <match_properties>
```

❖ Example

Refresh all statistics on virtual table `HANA1_T1`.

```
REFRESH STATISTICS ON HANA1_T1
```

Using linked database, refresh only `SIMPLE` statistics on table `T1` using remote source `HANA1` and schema `myschema`.

```
REFRESH STATISTICS ON HANA1.myschema.T1 TYPE SIMPLE
```

Related Information

[Monitor Details Using SAP HANA Cockpit \[page 1937\]](#)

[Monitor Details Using SAP HANA Studio \[page 1935\]](#)

13.1.19.4 Drop Statistics on a Virtual Table or Linked Database

Drop data statistic virtual objects that the query optimizer uses to make better decisions for query plans.

Prerequisites

One of the following:

- You created the virtual table you are dropping statistics from.
- You have the `ALTER` privilege on the object you are dropping statistics from.
- For linked database, you require the `LINKED DATABASE` object level privilege on the remote source, regardless of who created the remote source.

Procedure

1. To drop statistics on a virtual table, execute:

```
DROP STATISTICS ON [<schema_name>.] <virtual_table_name>  
[ [HAVING] <match_properties> ]
```

2. Using linked database, to drop statistics on a remote table, execute:

```
DROP STATISTICS ON <database_name>.<schema_name>.<remote_table_name>  
[ [HAVING] <match_properties> ]
```

❖ Example

The following statement drops all statistics on virtual table `HANA1_T1`.

```
DROP STATISTICS ON HANA1_T1
```

Using linked database, the following statement drops SIMPLE statistics on table `T1` using remote source `HANA1` and schema `myschema`.

```
DROP STATISTICS ON HANA1.myschema.T1 TYPE SIMPLE
```

Related Information

[Monitor Details Using SAP HANA Cockpit \[page 1937\]](#)

[Monitor Details Using SAP HANA Studio \[page 1935\]](#)

13.1.19.5 Retrieve Statistics from a Remote Source

On supported remote sources, statistics for virtual tables and linked database are retrieved by querying a remote table.

The name of the virtual table that is queried for statistics retrieval is:

Remote Source	Virtual Table Name
SAP HANA	SYSTEM.SDA_STATISTICS
SAP IQ	SYS_STATISTICS
Teradata	SYS_STATISTICS

`SYS_STATISTICS` is located on the default schema of the connection used to create the virtual tables.

When SIMPLE statistics are computed for a virtual table, the remote statistics table is queried first. If this table is not available (or has a different format), the standard behavior used to obtain statistics from remote sources is triggered, that is, statistics queries are sent for each column in order for the statistics to be remotely computed.

The schema of the remote statistics table is as follows:

Index	Name	Type	Precision	Description
1	SCHEMA_NAME	VARCHAR	128	Schema name
2	TABLE_NAME	VARCHAR	128	Table name
3	COLUMN_NAME	VARCHAR	128	Column name

Index	Name	Type	Precision	Description
4	MIN	VARCHAR	128	String representation of the min value
5	MAX	VARCHAR		String representation of the max value
6	COUNT_STAR	INTEGER		Count (*)
7	DCOUNT	INTEGER		Distinct count
8	COUNT	INTEGER		Count (used to count NULL values)

Virtual table statistics are stored in the DATA_STATISTICS system table.

13.1.20 Pool of Remote Connections

Use a pool of remote connections to enable multithreaded execution to be scaled out. A pool of remote connections can only be used for read-only remote sources (DML_MODE=readonly).

Provided no external updates (by third parties) occur on the remote source, full consistency is assured. However, if external updates occur in parallel on a declared read-only remote source, consistency can be ensured only within one SAP HANA thread.

Configuration Parameter for Pool of Remote Connections

Use the parameter `default_connections_pool_max_size` in the `smart_data_access` section of the `indexserver.ini` file to configure a pool of remote connections:

```
indexserver.ini/smart_data_access/default_connections_pool_max_size
```

- Default value: 3
- Highest value allowed: 50
- Value to disable the connection pool: 1

The maximum number of connections allowed in one pool is controlled by the value specified in `default_connections_pool_max_size`. Each SAP HANA connection has its own connection pool for each remote source it uses. The number of connections depends on the degree of multithreading of the executed statements, but cannot exceed the number specified in `default_connections_pool_max_size`. Also, each SAP HANA node has its own connection pool, so `default_connections_pool_max_size` applies per node and is not a global maximum. The query optimizer may decide to increase the degree of parallelism by using multiple SAP HANA nodes for the query execution.

13.1.20.1 Configure the Pool of Remote Connections Parameter

Configure the number of remote connections to enable for multi-threaded execution.

Prerequisites

Requires the INIFILE ADMIN system privilege.

Configure the Parameter Using SAP HANA Studio

Procedure

1. In the *Systems* view, connect to a system.
2. In the *SAP HANA Administration Console*, choose the *Configuration* tab.
3. Expand *indexserver.ini*, select *smart_data_access*, and choose *Add parameter* from the context menu.
4. Assign the values to *System*, enter a key *default_connections_pool_max_size*, and enter a value for the number of connections.
5. Choose *Finish* to create the parameter.

Configure the Parameter Using SQL

Procedure

In an SQL console, execute:

```
ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini','SYSTEM')
  SET ('smart_data_access', 'default_connections_pool_max_size')=<value>;
```

❖ Example

The following statement sets the number of connections to 10:

```
ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini','SYSTEM')
  SET ('smart_data_access', 'default_connections_pool_max_size')='10';
```

13.1.21 Results Caching for Virtual Tables and Linked Database

Only view results caching is supported for virtual tables and linked database.

You configure result caching using the `<cache_type>` parameter in the CREATE VIEW statement, or the `<alter_cache_settings_clause>` parameter in the ALTER VIEW statement.

Caching of user defined functions that reference virtual tables is not supported.

When querying a virtual table, if there is insufficient memory to cache the results, an internal error message appears, the query halts, and no results are cached. To prevent this scenario, modify the view definition if possible, to filter data before caching results, partition caching between multiple views, or disable result caching. You may need to clear the SQL plan cache after modifying the view definition. Using HINT IGNORE_PLAN_CACHE when executing a query allows the system to validate if there is sufficient memory to cache the results before execution and helps you decide if caching should be used. However, since remote data size is only an estimate, this may not prevent the error of insufficient memory.

When creating a view on virtual tables with caching configured, apply filters to the view definition not outside the view definition (part of the SQL query referencing the view) whenever possible. Using common filters within the view definition increases the usability and performances of the cache.

When results caching is enabled, only filters applied to the view definition are pushed down to the remote source. Filters outside of the view definition are not pushed. They are applied once the unfiltered results are cached. This potentially impacts how much data is retrieved from the remote source. When results caching is disabled, and the filter is outside the view, the filter may be pushed down to the remote source, returning less data, but the results are not cached. The next time the view is used, the new query is sent to the remote source and new data is returned.

Cached views can have a retention period, which when exceeded trigger a refresh of the view on next use. The frequency of refresh can impact performance. Regularly monitor cache use. The retention period of a view should reflect the nature of the data returned. Apply a longer retention period to views returning data more static in nature. We do not recommend that you cache views that are rarely used, regardless of the nature of the data. The constant refresh of the data could impact performance without any potential benefits from caching.

Data in a cached view is available to all users with SELECT privilege on the view. Remote privileges on virtual tables are validated only when a remote statement is executed, populating the results cache. The view results cache is available to all users. Any query can populate a view results cache, but the data cached is based on the remote credentials of the current user. As a result, users querying a cached view may see less data than expected if their remote privileges are greater than those of the user who populated the cache. Conversely, users with lesser remote privileges may see data in the shared cache to which they would not normally have access.

Use cached views only with remote sources using technical user credentials. For all other types of remote authentications, only SAP HANA side authorizations are enforced. When executing a query, to prevent populating the results cache with data specific to credentials, include the HINT RESULT_CACHE_NO_REFRESH clause.

Take careful consideration when enabling caching on views with virtual tables that point to data with static or dynamic analytical privileges or data that cannot be seen as a relational data set (that binds to relational operator rules). Remote constraints on data may not be enforced regardless of whether view results cache is enabled or not.

Result caching is disabled by default. When using HINT RESULT_CACHE in a query on views, if result cache is configured on the views, result caching is enabled for the query only.

Result caching is also controlled through the `<result_cache>` set of system configuration parameters within `indexserver.ini`.

See *SAP HANA SQL and System Views Reference* for full syntax to create, alter, and drop view cache, or to enable result caching.

Example

This statement creates a view called `view_tableA` on `tableA` with a static result cache retention period of 10 sec:

Sample Code

```
CREATE VIEW view_tableA AS SELECT * FROM tableA WITH STATIC CACHE RETENTION 10;
```

This statement drops the view changes the retention period from 10 to 5:

Sample Code

```
ALTER VIEW view_tableA ALTER STATIC CACHE RETENTION 5;
```

13.1.22 SAP HANA Automatic Failover Support

If a connection to a remote source becomes unavailable, the remote source automatically reconnects to the failover node.

Automatic failover functionality depends on the ODBC connection configuration and whether the remote source itself supports failover.

Currently, the following ODBC connections support failover:

- Hana
- SAP Adaptive Server Enterprise (ASE)
- Oracle

For details on enabling failover, see the specific database under [Creating a Remote Source \[page 1863\]](#).

13.1.23 Safe Mode for ODBC Connections

Provides the capability to load ODBC drivers and execute ODBC calls from within the scriptserver process. This reduces potential issues with the indexserver caused by third-party ODBC drivers.

Currently, smart data access loads the ODBC drivers required for communication with remote database as shared objects directly in the indexserver process. This can be problematic as most of these objects are third-party libraries and may not provide the stability and quality expected from SAP HANA. Any bug in the third-party ODBC driver library may cause the indexserver to crash, impacting customer productivity.

To address this, smart data access now allows the scriptserver to be used as a remote driver manager service. With this configuration the ODBC driver manager and ODBC drivers are loaded and executed in the scriptserver.

This functionality is configured in the `smart_data_access > odbc_adapters_in_scriptserver` property in the `indexserver.ini` file and can be enabled and disabled (default) for individual remote sources.

Enter a comma separated list of the adapter types to use the scriptserver. Valid entries include:

Supported Database	Drive Name
SAP IQ	iqodbc
SAP Adaptive Server Enterprise (ASE)	aseodbc
SAP HANA Accelerator for ASE	ets
SAP MaxDB	maxdb
Teradata	tdodbc
SQL Server	mssql
IBM DB2	db2
IBM Netezza	netezza
Oracle	oracle
Vora	voraodbc
Hive	hiveodbc
SAP IQ, ASE, HANA	hana_family
All supported databases excluding those included in hana_family	3rd_party
All supported databases	all

The scriptserver, disabled by default, must be enabled to use safe mode. See SAP Note [1650957](#) – SAP HANA Database: Starting the Script Server.

13.1.24 Setting Session Specific Information for Connections

Session specific client information can now be set for connections to HANA remote sources.

When creating remote sources, you add syntax to specify the session information as follows:

```
sessionVariable:<session_variable_name>=?
```

When the connection is made via smart data access to the remote source, the ? is replaced with the value of the variable in the local current session context and is added to the connection string. This allows values of local session variables to be used to set session variables on the remote HANA system.

This will be done for every connection being made to the remote source. Each connection to the local HANA system has its own session. When a session executes a query using a virtual table, it establishes a connection to the remote HANA system. The value of the variable used to generate the connection string depends on the value of the variable in the local HANA session.

For example, to set the session information on the remote HANA for the variables `APPLICATIONUSER` and `CDS_CLIENT`, if the value of the variables in the local session `APPLICATIONUSER` is `abc` and `CDS_CLIENT` is `100`. When the connection to the remote system is established, the following is added to the connection string:
`sessionVariable:APPLICATIONUSER=abc; sessionVariable:CDS_CLIENT=100.`

This feature is supported for HANA remote sources only. See [Create an SAP HANA Remote Source \[page 1863\]](#) for configuration information.

13.1.25 Smart Data Access System Parameters

Configuration parameters for smart data access are available in the `smart_data_access` and `linked_database` sections of the `indexserver.ini` file.

smart_data_access Section

Parameter Name	Description	Type	Length	Values	Default Value	Hidden
enable_remote_source_capability	Specifies the complexity of queries to be sent to the remote sources.	BOOLEAN		TRUE = any query in the remote source dialect can be sent for remote execution FALSE = only projections are sent for remote execution	TRUE	NO
linked_database_cleanup_interval	Specifies the interval in seconds to perform linked object house-keeping task	integer		Positive integer value in seconds 0 = task is disabled	0	NO

Parameter Name	Description	Type	Length	Values	Default Value	Hidden
semi_join_execution_strategies	Specifies the preferred order of semi-join execution strategies.	VARCHAR	16	IT = attempt of in-clause strategy followed by attempt of temporary table strategy TI, T = temporary table strategy I = in-clause strategy N = turns off the semi-join	IT	NO
semi_join_max_in_elements	Specifies maximum number of values in the IN clause for semi-join usage.	INTEGER		Positive integer value	1024	NO
semi_join_max_temp_table_cardinality	Maximum number of values to be inserted in a semi-join temp table.	INTEGER		Positive integer value	16384	NO
semi_join_min_temp_table_cardinality	Minimum number of values to be inserted in a semi-join temp table.	INTEGER		Positive integer value		YES
semi_join_reduction_factor	The estimated percentage reduction required for an attribute to be considered for semi-join reduction.	TINYINT		Positive integer value		YES
semi_join_virtual_table_threshold	Minimum number of estimated rows for fact subplan, to be considered for semi-join reduction.	TINYINT		Positive integer value		YES
virtual_table_format	Forces optimizer to use between column or row-based operators.	VARCHAR	16	ROW = row based COLUMN = column based AUTO = let the optimizer choose	ROW	NO

linked_database Section

Parameter Name	Description	Type	Length	Values	Default Value	Hidden
linked_data-base_cleanup_interval	Specifies the interval in seconds to perform linked object house-keeping task	integer		Positive integer value in seconds 0 = task is disabled	0	NO

13.1.26 Troubleshooting Smart Data Access

Find solutions to common smart data access problems.

If you don't find your issue in the troubleshooting topics, see [2352696 - SAP HANA Smart Data Access 2.0 Master Release Note](#) for a list of known issues.

13.1.26.1 Invalid Connection String Message When Querying a Google BigQuery Database

: When executing a query using Google BigQuery as a remote source, you receive a message regarding an invalid connection string.

Ensure that the DriverManagerEncoding property in the `simba.googlebigqueryodbc.ini` file on the SAP HANA host is configured to use UTF-16. See [Google BigQuery ODBC Driver \[page 1860\]](#).

13.2 SAP HANA Hadoop Integration

Regardless of structure, you can combine the in-memory processing power of SAP HANA with Hadoop's ability to store and process huge amounts of data.

SAP HANA is designed for high-speed data and analytic scenarios, while Hadoop is designed for very large, unstructured data scenarios. Hadoop can scale to thousands of nodes and is designed for use in large distributed clusters and to handle big data. Combining SAP HANA with Hadoop leverages Hadoop's lower storage cost and type flexibility with the high-speed in-memory processing power and highly structured data conformity of SAP HANA.

SAP HANA Hadoop integration is designed for users who may want to start using SAP HANA with their Hadoop ecosystem.

For information about SAP HANA Hadoop Integration and SAP HANA Spark controller installation, see [SAP HANA Hadoop Integration](#).

14 SAP HANA HDBSQL (Command-Line Reference)

SAP HANA HDBSQL is a command line tool for executing commands on SAP HANA databases.

Using SAP HANA HDBSQL, you can execute SQL statements and database procedures, as well as query information about the database and database objects. SAP HANA HDBSQL is installed with the SAP HANA software. It accesses databases both on your local computer and on remote computers.

Call SAP HANA HDBSQL with the command `hdbsql [options]` from the following location: `/usr/sap/<SID>/HDB<instance>/exe`. You can execute individual commands interactively or non-interactively. It is also possible to import commands from a file and execute them in the background.

14.1 SAP HANA HDBSQL Options

Set information for the database and for database objects by using SAP HANA HDBSQL commands.

In addition to SAP HANA HDBSQL commands, you can also enter an SQL statement or a database procedure. The statement or procedure must be enclosed in quotation marks.

Configuration Options

Use the following options to modify the operation of SAP HANA HDBSQL commands.

Database Session

Option	Description
<code>-attemptencrypt</code>	Specifies that encrypted data transmission is used. If the connection fails, then it attempts to use unencrypted connections.
<code>-d <database-name></code>	Specifies the name of the multitenant database container in a multiple-container system.
<code>-e</code>	Specifies that encrypted data transmission is used.
<code>-i <instance-number></code>	Specifies the instance number of the system.
<code>-n <host>[:<port>]</code>	Specifies the name of the computer on which the system is installed and optionally, the port number.
<code>-p <database-user-password></code>	Specifies the password for logging on to the database.
<code>-proxyhost <hostname></code>	Connects to a SOCKS5 proxy located at <code><hostname></code> .

Option	Description
<code>-proxypassword<pwd></code>	(Optional) Authenticates against a SOCKS5 proxy with <code><pwd></code> .
<code>-proxyport<port></code>	(Optional) Connects to a SOCKS5 proxy using <code><port></code> . The default is 1080.
<code>-proxyservicename<proxy-service-name></code>	(Optional) Specifies the proxy service name registered in Kerberos (for METHOD 01 authentication). The default is socks@proxyhost (in GSSAPI format). On Microsoft Windows, this option is mandatory, with the default being service/proxyhost@REALM (in SSPI format).
<code>-proxyuserid<userid></code>	(Optional) Authenticates against a SOCKS5 proxy using <code><userid></code> .
<code>-r</code>	Enforces the execution of SQL statements as statements rather than as prepared statements.
<code>-r</code>	Suppresses the use of prepared statements.
<code>-saml-assertion <file></code>	Uses a file to provide a SAML assertion.
<code>-S <sql-mode></code>	Specifies the SQL mode, either INTERNAL or SAPR3.
<code>-u <database-user></code>	Specifies the user name for logging on to the database.
<code>-U <user-store-key></code>	Uses credentials from the user store.
<code>-V <variable-definition>[...]<i>[prompt]</i> <i>[noprompt]</i></code>	<p>Specifies a variable. <code><variable-definition></code> can be one of the following:</p> <p>Ex- plicit Explicitly define the variable using <code><variable-name>=<variable-value></code>, with no spaces between the variable name, the equal sign (=), and the value.</p> <p>Posi- tio- nal Define the variable according to its position on the command line by using <code><variable-value></code>. The variable declaration's position on the command line is relative to other positional variables and determines which parameter it is replacing. For example, <code><value 1>,<value 2></code> means that &1 in the SQL script is replaced by <code><value 1></code>, and &2 is replaced by <code><value 2></code>.</p> <p>Spaces are only allowed if the entire list of definitions is quoted and special characters are escaped.</p> <p>Specify <i>[prompt]</i> to be prompted for a definition whenever an undefined variable is encountered. Specifying <i>[noprompt]</i> means that undefined variables are ignored and the command may fail if it contains undefined variables. The default is <i>[noprompt]</i>.</p>
<code>-z</code>	Switches off AUTOCOMMIT mode.

Input and Output

Option	Description
<code>-C <separator></code>	Specifies the separator used to separate individual commands when importing commands from a file. The default value is ;.
<code>-I <file></code>	Imports commands from a batch file.
<code>-m</code>	Activates multiple-line mode for entering SAP HANA HDBSQL commands.
<code>-O <file></code>	Writes the results to a file.
<code>-x</code>	Suppresses additional output, such as the number of selected rows in a result set.
<code>-resultencoding <encoding></code>	Forces output encoding for result data. Can be one of UTF8 , LATIN1 , or AUTO (the default).
<code>-qto</code> <code>-querytimeout</code>	Sets a server-side timeout for all SQL operations, in seconds. If any SAP HANA HDBSQL SQL operation exceeds the maximum timeout value on the server, then it is canceled with a server error message. Setting a timeout value of 0 disables the timeout (the default).
<code>-quiet</code>	Hides the SAP HANA HDBSQL welcome banner.
<code>-separatorownline</code>	Deprecated, do not use. Nested BEGIN...END blocks are supported by default without modifying the input.
<code>-strictSeparatorLine</code>	Removes the parsing of a single quote, double quote, and BEGIN...END nesting. Separator line matching is strict and no leading or trailing spaces are allowed. When <code>-c</code> is not used, the default separator is a semi-colon on its own line. Batch tracing (<code>-f</code> or <code>-fn</code>) also produces output for comment-only batches even when they are not sent to the server.
<code>-br <reset-command></code>	Specifies a reset command (for example, <code>reset</code>), which tells SAP HANA HDBSQL to ignore the most recent query sent to the server. To specify <code>batchreset</code> , you must also set the <code>separatorownline</code> option. The reset command should not be an SQL statement or a separator. Setting <code>batchreset</code> on the same line as the separator command results in an error because it is not a valid query.

Option	Description
<code>-printoutput { OFF MESSAGE }</code>	<p>Specifies how to handle the SQLSCRIPT_PRINT library output. The parameters are case sensitive.</p> <p>MES- SAGE This is the default setting if you do not specify <code>-printoutput</code>. MESSAGE causes the SQLSCRIPT_PRINT library output to be printed on standard output. The <code>-x</code> option suppresses the SQLSCRIPT_PRINT library output even if <code>-printoutput MESSAGE</code> is set.</p> <p>OFF OFF suppresses output from the SQLSCRIPT_PRINT library when you call a stored procedure.</p>

<code>-history <#></code>	Specifies the number of items to keep in the history buffer. The default is 50.
---------------------------------	---

Formatting Output

Option	Description
<code>-A</code>	Returns the result set in an aligned format.
<code>-a</code>	Suppresses the output of the column names in the result set.
<code>-C</code>	Suppresses escape output format.
<code>-b <maximum-length> all</code>	Defines the maximum number of characters for the output of LOB values (the default value is 32 bytes). Specifying <code>-b all</code> displays the whole binary length.
<code>-f</code>	Returns all SQL statements that are sent to the database instance.
<code>-fn</code>	Returns all SQL statements that are sent to the database instance and formats them with numbered lines. Numbered lines make it easier to determine on which file line a potential error has occurred.
<code>-F <separator></code>	Specifies which string SAP HANA HDBSQL uses as a separator between the individual columns of the result set. The default value is .
<code>-g <>null-value></code>	Specifies the character for NULL values in the result set. The default value is ?.
<code>-p <prefix></code>	Use <code><prefix></code> as the row prefix for printing. The default value is .
<code>-q <suffix></code>	Use <code><suffix></code> as the row suffix for printing. The default value is .

Option	Description
<code>-oldexectimes</code>	Uses SAP HANA 1.0 execution-only timing. SAP HANA HDBSQL in SAP HANA 1.x only reports time for client and server executions, not fetches for result sets. As of SAP HANA 2.0, SAP HANA HDBSQL includes times for executions and fetches by default.
<code>-Q</code>	Outputs each column of the result set in a new row.
<code>-j</code>	Switches off the page by page scroll output.

i Note

By default, SAP HANA HDBSQL removes whitespace of the end of column data values. To prevent this, add `-Z CHOPBLANKS=0` to the SAP HANA HDBSQL command line.

Other

Option	Description
<code>-h</code>	Displays the help.
<code>-t</code>	Outputs debug information.
<code>-T <file></code>	Activates the SQLDBC trace, which writes the trace data to the specified file.
<code>-v</code>	Displays version information about the SAP HANA HDBSQL program.

SSL Options

Option	Description
<code>-sslprovider <provider></code>	Specifies the cryptographic service provider that is used for SSL connections (one of commoncrypto, sapcrypto, or mscrypto).
<code>-sslkeystore <key-store-file></code>	Specifies the SSL keystore name.
<code>-ssltruststore <trust-store-file></code>	Specifies the SSL truststore name.
<code>-ssltrustcert <certificate-file></code>	Skips certificate validation.
<code>-sslhostnameincert <hostname></code>	Specifies the hostname of the server for which the certificate has been granted.
<code>-sslcreatecert</code>	Creates a self-signed certificate.

Interactive Options

Use the following options when operating SAP HANA HDBSQL in interactive mode.

Command	Description
<code>\?</code>	Displays all HDBSQL commands.
<code>\h[elp]</code>	
<code>\a[utocommit] [ON OFF]</code>	Switches AUTOCOMMIT mode on or off.
<code>\a[ign] [ON OFF]</code>	Controls whether SQL statement results are formatted.
<code>\e[scape] [ON OFF]</code>	Switches the escape output format on or off.
<code>\c[onnect]</code>	Logs a user onto the database.
<code>\dc [<pattern>]</code>	<p>Lists all table columns that correspond to the specified [<pattern>] and to which the current user has access.</p> <p>[<pattern>] is specified as follows: [<schema>.] [<object-name>]. The following placeholders are supported:</p> <ul style="list-style-type: none"> • For one character: _ • For any number of characters: % <p>If a pattern is not specified, then the system returns information about all table columns to which the current user has access.</p> <p>This command returns the following information:</p> <ul style="list-style-type: none"> • Column name • Data type • Column length • Null value permitted or not • Position of the column in primary key of table (if applicable)
<code>\de [<pattern>]</code>	<p>Lists all the indexes of database objects that correspond to the specified [<pattern>].</p> <p>[<pattern>] is specified as follows: [<schema>.] [<object-name>]. The following placeholders are supported:</p> <ul style="list-style-type: none"> • For one character: _ • For any number of characters: % <p>If a pattern is not specified, then the system returns information about all indexes for database objects to which the current user has access.</p> <p>This command returns the following information:</p> <ul style="list-style-type: none"> • Index name • Columns contained in the index • Position of column in the index • Specifies whether the index is UNIQUE • Sort sequence
<code>\di[sconnect]</code>	Logs the user off of the database.

Command	Description
<code>\dp [<pattern>]</code>	<p>Lists all database procedures that correspond to the specified [<pattern>].</p> <p>[<pattern>] is specified as follows: [<schema>.] [<object-name>]. The following placeholders are supported:</p> <ul style="list-style-type: none"> • For one character: _ • For any number of characters: % <p>If a pattern is not specified, then the system returns information about all database procedures to which the current user has access.</p> <p>This command returns the following information:</p> <ul style="list-style-type: none"> • Schema name • Name of the database procedure • Package to which the database procedure is assigned
<code>\ds [<name>]</code>	<p>Lists all schemas that correspond to the specified [<name>].</p> <p>[<name>] is specified as follows: [<schema>.] [<object-name>]. The following placeholders are supported:</p> <ul style="list-style-type: none"> • For one character: _ • For any number of characters: % <p>If a pattern is not specified, then the system returns information about all schemas to which the current user has access.</p> <p>This command returns the following information:</p> <ul style="list-style-type: none"> • Schema Name • Owner
<code>\dt [<pattern>]</code>	<p>Lists all tables that correspond to the specified [<pattern>].</p> <p>[<pattern>] is specified as follows: [<schema>.] [<object-name>]. The following placeholders are supported:</p> <ul style="list-style-type: none"> • For one character: _ • For any number of characters: % <p>If a pattern is not specified, then the system returns information about all tables to which the current user has access.</p> <p>This command returns the following information:</p> <ul style="list-style-type: none"> • Schema name • Table name • Table type

Command	Description
<code>\du [<name>]</code>	<p>Lists all database users that correspond to the specified [<name>].</p> <p>[<name>] is specified as follows: [<schema>.] [<object-name>]. The following placeholders are supported:</p> <ul style="list-style-type: none"> • For one character: _ • For any number of characters: % <p>If a name is not specified, then the system returns information about all database users to which the current user has access.</p> <p>This command returns the following information:</p> <ul style="list-style-type: none"> • Name of the database user • User properties
<code>\dv [<pattern>]</code>	<p>Lists all views that correspond to the specified [<pattern>].</p> <p>[<pattern>] is specified as follows: [<schema>.] [<object-name>]. The following placeholders are supported:</p> <ul style="list-style-type: none"> • For one character: _ • For any number of characters: % <p>If a pattern is not specified, then the system returns information about all views to which the current user has access.</p> <p>This command returns the following information:</p> <ul style="list-style-type: none"> • Schema name • View name • View types
<code>\edit [<file>]</code>	Writes the command buffer to the specified file where you can edit it with an editor.
<code>\f[ieldsep] <separator></code>	Uses the specified separator character to separate the individual fields of the result. The default is ,.
<code>\g</code>	Executes the commands in the command buffer and returns the results.
<code>\h[istory] <#></code>	Specifies the number of items to keep in the history buffer. The default is 50.
<code>\i[nput] <file></code>	Imports commands from the specified batch file.
<code>\m[ode] { INTERNAL SAPR3 }</code>	Changes the SQL mode.
<code>\mu[ltiline] { ON OFF }</code>	Switches multiple line mode on or off.
<code>\o[utput] <file></code>	Redirects the result to a file.
<code>\pa[ger]</code>	Displays results consecutively (not page by page).
<code>\p[rint]</code>	Displays the current command buffer.
<code>\qto \querytimeout</code>	Sets a server-side timeout for all SQL operations, in seconds. If any SAP HANA HDBSQL SQL operation exceeds the maximum timeout value on the server, then it is canceled with a server error message. Setting a timeout value of 0 disables the timeout. This is the default.
<code>\q[uit]</code>	Exits HDBSQL.

Command	Description
<code>\r[eset]</code>	Deletes the current command buffer.
<code>\read <file></code>	Reads commands from the specified batch file.
<code>\ro[wsep] <separator></code>	Uses the specified separator character to separate the individual rows of the result.
<code>\s[tatus]</code>	Displays general information about the database.
<code>\v[-]</code>	Lists all currently defined variables.
<code>\vd <variable-name> <value></code>	Defines a <code><variable-name></code> to be replaced by <code><value></code> , or replaces the value stored with a new value if <code><variable-name></code> already is defined.
<code>\vu <variable-name></code>	Undefines the specified variable.
<code>\vc</code>	Clears all of the currently defined variables.
<code>\vs { ON OFF }</code>	Switches variable substitution on or off.
<code>\ve { ON OFF }</code>	Switches variable escaping on or off.
<code>\vp { ON OFF }</code>	Switches variable prompting on or off.

14.2 Log On to a Database

Log on to the database as a database user to use SAP HANA HDBSQL interactively and to execute commands.

Prerequisites

The user logging on must be a database user. If you do not specify the user name and password of a database user, then the logon is attempted using Kerberos authentication.

Procedure

- Log onto a database using SAP HANA HDBSQL with either a one-step or two-step process.
 - a. To log onto a database in one step, with a user name and password, run one of the following commands:

Option	Action
Log onto a database in a single-container system	Run the following command all on one line: <pre>hdbsql -n <host> -i <instance> -u <database_user> - p <database_user_password></pre>
Log onto a database in a multitenant database container	Run the following command all on one line: <pre>hdbsql -n <host> -i <instance> -u <database_user> - p <database_user_password> -d <database_name></pre>

- b. To log onto a database in two steps, with a user name and password, run the following commands:
1. Start SAP HANA HDBSQL by running `hdbsql`.
 2. Log on to the database by running on of the following commands:

Option	Action
Log onto a database in a single-container system	Run the following command all on one line: <pre>\c -n <host> -i <instance> -u <database_user> - p <database_user_password></pre>
Log onto a database in a multitenant database container	Run the following command all on one line: <pre>\c -n <host> -i <instance> -u <database_user> - p <database_user_password> -d <database_name></pre>

i Note

You can log on with user credentials for the secure user store (`hdbuserstore`) with `-U <user_key>`. For more information, see *Secure User Store (hdbuserstore)* in the *SAP HANA Security Guide*.

Results

The user is connected to the system or the multitenant database container.

❖ Example

For one-step logon to the system on the PARMA host with instance number 01 as database user MONA with the password RED, run the following command:

```
hdbsql -n PARMA -i 1 -u MONA -p RED
```

For one-step logon to the system database of system MDB1 on MYHOST with instance number 2 as database user SYSTEM with password BLUE, run the following command:

```
hdbsql -n MYHOST -i 2 -u SYSTEM -p BLUE -d SYSTEMDB
```

14.3 Run Commands

Run SAP HANA HDBSQL commands in interactive and non-interactive mode.

Prerequisites

You must be logged on to the database.

Context

To execute an SQL statement or a database procedure as a command, place the statement or procedure in quotation marks.

Procedure

- Run a command in interactive (session) mode as follows:
 - a. Call SAP HANA HDBSQL by running the following command: `hdbsql`
 - b. Type in the command and press **Enter**.
SAP HANA HDBSQL runs the command.
 - c. Exit SAP HANA HDBSQL by running one of the following commands: `exit` | `quit` | `\q`
- Run a command in non-interactive (command) mode as follows:

```
hdbsql [options] <command>
```

SAP HANA HDBSQL runs the command and then exits.

- Run multiple commands from a batch file as follows:

```
hdbsql [<options>] -I <file>
```

SAP HANA HDBSQL imports the commands from the specified file and processes them in the background. Specify the separator used in the batch file to separate individual commands by using the `-c <separator>` command line option. The default value is a semicolon (;).

i Note

If you run commands from a batch file, then AUTOCOMMIT mode is activated by default. If you deactivate AUTOCOMMIT mode, then the batch file must contain an explicit COMMIT statement to ensure that SAP HANA HDBSQL executes the SQL statements immediately after the batch file has been imported.

Example

Run the following command to display general information about the database in command mode with simultaneous database logon:

```
hdbsql -n localhost -i 1 -u USER1 -p Password123 \s
```

The above command returns the following result:

```
host: wdfd00245293a:30015
database: ORG
user: USER1
kernel version: 1.00.38.368649
SQLDBC version: libSQLDBC_HDB 1.00.38.368649 Build 0000000-0120
autocommit: ON
```

Run the following command all on one line to execute the SELECT statement in command mode with simultaneous database logon:

```
hdbsql -n localhost -i 1 -u USER1 -p Password123
"SELECT CNO,TITLE,FIRSTNAME,NAME,ZIP FROM HOTEL.CUSTOMER"
```

The above command returns the following result:

```
CNO | TITLE | FIRSTNAME | NAME | ZIP
-----+-----+-----+-----+-----
3000 | Mrs | Jenny | Porter | 10580
3100 | Mr | Peter | Brown | 48226
3200 | Company | ? | Datasoft | 90018
3300 | Mrs | Rose | Brian | 75243
3400 | Mrs | Mary | Griffith | 20005
3500 | Mr | Martin | Randolph | 60615
3600 | Mrs | Sally | Smith | 75243
3700 | Mr | Mike | Jackson | 45211
3800 | Mrs | Rita | Doe | 97213
3900 | Mr | George | Howe | 75243
4000 | Mr | Frank | Miller | 95054
4100 | Mrs | Susan | Baker | 90018
4200 | Mr | Joseph | Peters | 92714
4300 | Company | ? | TOOLware | 20019
4400 | Mr | Antony | Jenkins | 20903
(15 rows selected) * Ok
```

Run multiple commands imported from a batch file in command mode:

```
hdbsql [<options>] -I CITES
```

The file contains the following statements for execution:

```
CREATE TABLE city
(zip NCHAR (5) PRIMARY KEY,
name NCHAR(20),
state NCHAR(2) );
CREATE TABLE customer
(cno INTEGER PRIMARY KEY,
title NCHAR (7),
firstname NCHAR (10),
name NCHAR (10),
zip NCHAR (5),
address NCHAR (25));
\dt customer;
COMMIT
```

14.4 Run Long Commands in Multiple-Line Mode

Multiple-line mode enables you to enter long commands, for example a long SQL statement on several lines. SAP HANA HDBSQL stores multiple-line commands in an internal command buffer.

Prerequisites

To run some commands, you must be logged on to the database.

Procedure

1. Activate multiple-line mode by running one of the following commands:
 - Call option: `hdbsql [<options>] -m`
 - SAP HANA HDBSQL command: `\mu ON`
2. Enter the command.
To start a new line, press **Enter**.
3. Run the command in one of the following ways:
 - Close the last line of the command by entering a semicolon and pressing **Enter**.
 - SAP HANA HDBSQL command: `\g`.

❁ Example

1. Log onto the SAP HANA database as user MONA with the password RED by running the following command: `hdbsql -n localhost -i 1 -u MONA,RED`

2. Activate multiple line mode by running the following command: `\mu ON`
3. Enter a multiple-line SQL statement:

```
SELECT ROUND(SUM("M")/1024/1024/1024,2) AS "Peak Used Memory GB" FROM
(SELECT SUM(CODE_SIZE+SHARED_MEMORY_ALLOCATED_SIZE) AS "M" FROM
SYS.M_SERVICE_MEMORY UNION SELECT SUM(INCLUSIVE_PEAK_ALLOCATION_SIZE) AS
"M" FROM M_HEAP_MEMORY_RESET WHERE DEPTH = 0)
```

4. Execute the SQL statement by entering the following command: `\g`

14.5 Edit Long Commands in an External File

If you have entered a long command in SAP HANA HDBSQL in multiple-line mode, then you can change it later by editing the command buffer in an external file and then re-running it.

Prerequisites

You have already run the command.

Procedure

1. To export the contents of the command buffer to an external file, run the following command:

```
\e <[file]>
```

You must enter the complete file path and file name. If you do not specify a file, then SAP HANA HDBSQL generates a temporary file.

The system opens the file in an editor. To determine which editor is used, SAP HANA HDBSQL evaluates the environment variables `HDBSQL_EDITOR`, `EDITOR`, and `VISUAL` in succession. If you have not set any of these environment variables, then the visual editor is used on Linux and UNIX. For more information about setting environment variables, see your operating system documentation.

2. Make the required changes to the file.
3. Save the file in the editor and then close the file and the editor.

Results

You have changed the contents of the command buffer and can now execute the changed command by running the command `\g`.

14.6 Redirect Results to a File

Redirect the result of one or more SAP HANA HDBSQL commands to a file.

Prerequisites

To redirect results to a file, you must be logged on to the database.

Procedure

1. Run the following command:

```
\o <file>
```

You must enter the full path of the file.

2. Run the command whose result is being redirected to the file.
To run multiple commands in succession, press **Enter** after each command.
3. To stop redirection to a file, run the following command: \o.

❖ Example

To export a list of all schemas and all entries in the table HOTEL.CUSTOMER to an external file, perform the following steps.

1. Log onto the SAP HANA database as user MONA with the password RED by running the following command:

```
hdbsql -n localhost -i 1 -u MONA, RED
```

2. Create the file `c:\tmp\redirected.txt` then redirect the command result(s) to this file by running the following command:

```
\o c:\tmp\redirected.txt
```

3. Request information about all schemas by running the following command: \ds
4. Select all rows in the table HOTEL.CUSTOMER by executing the following statement:

```
SELECT * FROM HOTEL.CUSTOMER
```

5. Stop redirection to the file by running \o.

The `redirected.txt` file now contains the following content:

```
| Schema | Owner name |
| ----- | - |
| MDX_TE | SYSTEM |
| SECURI | SECURITY1 |
| SOP_PL | SYSTEM |
| SYS | SYS |
```

```

| SYSTEM | SYSTEM |
| _SYS_B | _SYS_REPO |
| _SYS_B | _SYS_REPO |
| _SYS_R | _SYS_REPO |
| _SYS_S | _SYS_STATISTICS |
| CNO | TITLE | FIRSTNAME | NAME | ZIP | ADDRESS|
| ----- | ----- | ----- | ----- | ----- | -----|
| 3200 | Company | ? | Datasoft | 90018 | 486 Maple Str.|
| 3400 | Mrs | Mary | Griffith | 20005 | 3401 Elder Lane|
| 3500 | Mr | Martin | Randolph | 60615 | 340 MAIN STREET, #7|
| 3600 | Mrs | Sally | Smith | 75243 | 250 Curtis Street|
| 3700 | Mr | Mike | Jackson | 45211 | 133 BROADWAY APT. 1|
| 3900 | Mr | George | Howe | 75243 | 111 B Parkway, #23|
| 4000 | Mr | Frank | Miller | 95054 | 27 5th Str., 76|
| 4400 | Mr | Antony | Jenkins | 20903 | 55 A Parkway, #15|

```

14.7 Substitution Variables

Substitution variables are processed by SAP HANA HDBSQL before the query that includes them is sent to the server.

Identifying Variables

Substitution variables are identified by either the prefix `&` or the prefix `&&`. If a variable is prefixed by `&`, then you are prompted to define the variable. If the variable is prefixed by `&&`, then you are prompted to specify a new value only if the variable has not already been defined.

You can also use the `[prompt]` and `[noprompt]` parameters of the `-V` option to control whether SAP HANA HDBSQL prompts you to define variables that it encounters in the script.

Rules for Defining Variables on the Command Line

Quoting and Spacing

Variable definitions can be non-quoted or quoted. Variable values that contain spaces must be single and double-quoted on the command line, and single-quoted in a SQL script file.

Single quotes on the command line must be escaped.

Non-quoted variable definitions cannot have spaces between the equal sign (`=`) and the variable name or definition, or between a comma and a variable definition.

The following example is not quoted because the variable definition contains no spaces and there are no spaces between the equal sign and the variable name and variable definition.

```
<variable-name>=<variable-value-with-no-spaces>
```

The following example is single and double-quoted because the variable definition contains spaces and is defined on the command line.

```
"<variable-name> = '<variable-value-with-spaces>'"
```

The following example is single-quoted only because the variable value contains spaces and is defined in a SQL script file.

```
<variable-name> = '<variable-value-with-spaces>'
```

Escape Characters

The following characters: a-z, A-Z, 0-9, |, and _ are allowed in variable names.

On the command line, variable values that contain quotes and backslashes must be escaped by using a backslash (\) and the string must be quoted.

The following example shows a variable value defined on the command line with special characters escaped, single quotes so that SAP HANA HDBSQL can process a variable value that includes spaces, and double quotes for the command shell to process a variable value that includes spaces.

```
-V VAR_1=\"\"Sample string with\\ , and backslash \\\\\"\"'
```

The variable definition above is parsed as follows:

```
'\"Sample string with\\ , and backslash \\\\"'
```

Important Disclaimer for Features in SAP HANA Platform

For information about the capabilities available for your license and installation scenario, refer to the Feature Scope Description (FSD) for your specific SAP HANA version on the [SAP HANA Platform webpage](#).

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.

About the icons:

- Links with the icon : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
 - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
 - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon : You are leaving the documentation for that particular SAP product or service and are entering a SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.

The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Gender-Related Language

We try not to use gender-specific word forms and formulations. As appropriate for context and readability, SAP may use masculine word forms to refer to all genders.

© 2018 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.