

PUBLIC

SAP HANA Platform SPS 12  
Document Version: 1.2 – 2018-01-24

# SAP HANA Administration Guide

The SAP logo is located in the bottom left corner of the page. It consists of the letters 'SAP' in a bold, white, sans-serif font, set against a dark blue rectangular background. The background of the entire page is a long-exposure photograph of a complex highway interchange at night, with light trails from cars and streetlights creating a vibrant, blue-toned scene.

---

# Content

- 1 Getting Started. . . . . 8**
  
- 2 Overview of SAP HANA Architecture. . . . . 10**
  - 2.1 Multiple-Host Systems. . . . . 14
  - 2.2 Multitenant Database Containers. . . . . 15
    - Server Architecture of Multiple-Container Systems. . . . . 15
    - Scale-Out Architecture of Multiple-Container Systems. . . . . 17
    - The System Database. . . . . 18
    - Administration of Multitenant Database Containers. . . . . 19
  
- 3 SAP HANA Administration Tools. . . . . 21**
  - 3.1 SAP HANA Cockpit. . . . . 22
    - Open SAP HANA Cockpit. . . . . 23
    - Open SAP HANA Cockpit from SAP HANA Studio. . . . . 25
    - Customize the Homepage of SAP HANA Cockpit. . . . . 27
    - Configure Access to Content in SAP HANA Cockpit. . . . . 31
    - Tile Catalogs for Administration and Monitoring. . . . . 33
    - Tile Catalogs for Security Administration. . . . . 45
    - Tile Catalog for Platform Lifecycle Management. . . . . 52
    - Roles for Tile Catalogs. . . . . 54
    - Roles Granted to Database User SYSTEM. . . . . 56
    - SAP HANA Cockpit for Offline Administration. . . . . 56
  - 3.2 SAP HANA Studio. . . . . 60
    - Open the SAP HANA Administration Console. . . . . 61
    - Execute SQL Statements in SAP HANA Studio. . . . . 65
    - Managing SAP HANA Systems in SAP HANA Studio. . . . . 67
    - SAP HANA Studio Administration Preferences. . . . . 86
  
- 4 System Administration. . . . . 91**
  - 4.1 Starting and Stopping SAP HANA Systems. . . . . 91
    - Starting and Stopping Systems in SAP HANA Cockpit. . . . . 91
    - Starting and Stopping Systems in SAP HANA Studio. . . . . 97
  - 4.2 Managing Multitenant Database Containers. . . . . 104
    - Creating and Configuring Tenant Databases. . . . . 104
    - Monitoring and Managing Tenant Databases. . . . . 129
    - Managing Resources in Multiple-Container Systems. . . . . 143
    - Copying and Moving Tenant Databases Between Systems. . . . . 145
    - Using SAP Web Dispatcher for Load Balancing with Tenant Databases. . . . . 182

	Tutorial: Migrating SAP DB Control Center to a Tenant Database. . . . .	196
4.3	Configuring SAP HANA System Properties (INI Files). . . . .	212
	Configuration Parameters in Multiple-Container Systems. . . . .	213
	Change a System Property. . . . .	217
	Reset a System Property. . . . .	219
	Reserve Connections for Administrators. . . . .	220
	Configure System Usage Type. . . . .	221
4.4	Managing SAP HANA Licenses. . . . .	222
	License Keys. . . . .	222
	Check the Current License Key. . . . .	224
	Install a Permanent License. . . . .	224
	Delete an Existing Permanent License Key. . . . .	226
4.5	Monitoring the SAP HANA Database. . . . .	227
	Monitoring in SAP HANA Studio. . . . .	227
	Monitoring in SAP HANA Cockpit. . . . .	290
4.6	Managing Tables. . . . .	331
	Table Types in SAP HANA. . . . .	332
	Basic Table Management in SAP HANA Studio. . . . .	335
	Table and Catalog Consistency Checks. . . . .	348
	Memory Management in the Column Store. . . . .	352
	The Delta Merge Operation. . . . .	356
	Data Compression in the Column Store. . . . .	368
	Table Partitioning. . . . .	372
	Table Replication. . . . .	398
	Table Placement. . . . .	414
4.7	Workload Management. . . . .	418
	Workload in the Context of SAP HANA. . . . .	418
	Controlling CPU Consumption. . . . .	424
	Controlling Parallel Execution of SQL Statements. . . . .	427
	Setting a Memory Limit for SQL Statements. . . . .	431
	Managing Workload with Workload Classes. . . . .	435
	Example Workload Management Scenarios. . . . .	444
4.8	Scheduling of Recurring Administration Tasks. . . . .	447
4.9	Hardware Checks for Tailored Data Center Integration. . . . .	447
	Install the SAP HANA Hardware Configuration Check Tool. . . . .	448
4.10	SAP Solution Manager for SAP HANA. . . . .	449
	SAP Solution Manager for SAP HANA Administration. . . . .	450
	Configuring an SAP HANA System to Connect to the System Landscape Directory (SLD). . . . .	451
	Change the Default SLD Data Supplier Configuration. . . . .	457
4.11	Getting Support. . . . .	460
	Diagnosis Files. . . . .	460

	Configure Traces in SAP HANA Studio. . . . .	465
	Configure Trace File Rotation. . . . .	480
	Troubleshooting an Inaccessible or Unresponsive SAP HANA System . . . . .	481
	Problem Analysis Using hdbcons. . . . .	484
	Collecting Diagnosis Information for SAP Support. . . . .	484
	Collecting Performance Monitor Data for SAP Support. . . . .	496
	Open a Support Connection. . . . .	498
<b>5</b>	<b>SAP HANA Lifecycle Management. . . . .</b>	<b>499</b>
5.1	SAP HANA Platform Lifecycle Management. . . . .	500
	About the SAP HANA Database Lifecycle Manager (HDBLCM). . . . .	502
	Configuring the SAP HANA System. . . . .	531
	Changing the SAP HANA System. . . . .	583
	Managing SAP HANA System Components. . . . .	596
	Check the Installation Using the Command-Line Interface. . . . .	597
5.2	SAP HANA Application Lifecycle Management. . . . .	599
	Installing and Updating SAP HANA Products and Software Components. . . . .	600
	Installing and Updating Products and Software Components in SAP HANA XS Advanced Model . . . . .	607
	Configuring SAP HANA Applications with the Process Engine. . . . .	625
5.3	SAP HANA Content. . . . .	630
	SAP HANA Archive Types. . . . .	630
	Deploy a Product Archive (*.ZIP). . . . .	631
	Deploy a Delivery Unit Archive (*.tgz). . . . .	632
<b>6</b>	<b>Security Administration. . . . .</b>	<b>633</b>
6.1	Monitoring Critical Security Settings in SAP HANA Cockpit. . . . .	633
	View Status of Security Settings. . . . .	634
	Network Security Details. . . . .	635
6.2	Managing SAP HANA Users. . . . .	637
	Database Users. . . . .	638
	Operating System User <sid>adm. . . . .	647
	User Authentication and Single-Sign On. . . . .	648
	User Authorization. . . . .	669
	Provisioning Users. . . . .	702
6.3	Auditing Activity in SAP HANA Systems. . . . .	718
	Managing Auditing in the SAP HANA Cockpit. . . . .	719
	Managing Auditing in the SAP HANA Studio. . . . .	726
	Audit Trail Targets. . . . .	731
	Best Practices and Recommendations for Creating Audit Policies. . . . .	732
6.4	Managing Encryption Keys. . . . .	734
	Server-Side Data Encryption. . . . .	735

	Client-Side Data Encryption (hdbuserstore) . . . . .	748
6.5	Managing Encryption of Data Volumes in the SAP HANA Database. . . . .	749
	Enable Data Volume Encryption in an Existing SAP HANA System. . . . .	750
	Change the Page Encryption Key. . . . .	755
	Disable Data Volume Encryption. . . . .	756
	Data Volume Encryption in Multitenant Database Containers. . . . .	757
6.6	Managing Client Certificates in the SAP HANA Database. . . . .	758
	Client Certificates. . . . .	760
	Certificate Collections . . . . .	762
	View Certificates in the Certificate Store. . . . .	764
	View Certificate Collections. . . . .	765
	Import a Trusted Certificate into the Certificate Store. . . . .	766
	Create a Certificate Collection. . . . .	767
	Set the Purpose of a Certificate Collection. . . . .	768
	SQL Statements and Authorization for In-Database Certificate Management. . . . .	770
<b>7</b>	<b>Availability and Scalability . . . . .</b>	<b>773</b>
7.1	High Availability Navigation Support. . . . .	773
7.2	High Availability for SAP HANA. . . . .	774
	SAP HANA Disaster Recovery Support. . . . .	776
	SAP HANA Fault Recovery Support. . . . .	779
	Setting Up System Replication. . . . .	782
	Setting Up Host Auto-Failover. . . . .	850
	Implementing a HA/DR Provider. . . . .	854
7.3	SAP HANA Database Backup and Recovery. . . . .	868
	Savepoints and Redo Logs. . . . .	869
	Points to Note About Backup and Recovery. . . . .	870
	Authorizations for Backup and Recovery. . . . .	879
	SAP HANA Backup. . . . .	880
	SAP HANA Recovery. . . . .	936
	Copying a Database Using Backup and Recovery. . . . .	974
	Housekeeping for Backup Catalog and Backup Storage. . . . .	990
	Planning Your Backup and Recovery Strategy. . . . .	994
	Reference: Backup Alerts. . . . .	996
7.4	Scaling SAP HANA. . . . .	998
	Aspects of Scalability. . . . .	999
	Configuring the Network for Multiple Hosts. . . . .	1000
	Mapping Host Names for Database Client Access. . . . .	1002
	Add or Remove Hosts from an SAP HANA System. . . . .	1003
	Scaling SAP HANA Extended Application Services (XS). . . . .	1003
	Starting and Stopping Distributed SAP HANA Systems Using SAPControl. . . . .	1004
	Table Distribution in SAP HANA. . . . .	1004

	Monitor Table Distribution. . . . .	1006
	Redistribution of Tables in a Distributed SAP HANA System. . . . .	1007
<b>8</b>	<b>Maintaining the Application Services Run-Time Environment. . . . .</b>	<b>1017</b>
8.1	Maintaining the SAP HANA XS Classic Model Run Time. . . . .	1017
	SAP HANA XS Administration Tools. . . . .	1018
	SAP HANA XS Administration Roles. . . . .	1020
	SAP HANA XS Configuration Parameters. . . . .	1022
	Maintaining Application Runtime Configurations. . . . .	1029
	Managing Trust Relationships. . . . .	1043
	Maintaining SAML Providers. . . . .	1050
	Maintaining SMTP Server Configurations. . . . .	1059
	Maintaining HTTP Access to SAP HANA. . . . .	1066
	Maintaining Single Sign-On for SAP HANA XS Applications. . . . .	1074
	Maintaining User Self Service Tools. . . . .	1085
	Scheduling XS Jobs. . . . .	1110
	Maintaining Translation Text Strings. . . . .	1124
	Maintaining HTTP Traces for SAP HANA XS Applications. . . . .	1131
8.2	Maintaining the SAP HANA XS Advanced Model Run Time. . . . .	1139
	SAP HANA XS Advanced Administration Tools. . . . .	1140
	Role Collections for XS Advanced Administrators. . . . .	1141
	Monitoring the SAP HANA XS Advanced Model Runtime. . . . .	1143
	Maintaining Organizations and Spaces in SAP HANA XS Advanced Model. . . . .	1147
	Setting Up Security Artifacts. . . . .	1154
	Building Roles for SAP HANA XS Advanced Model Applications. . . . .	1155
	Managing SAML Identity Providers in XS Advanced. . . . .	1162
	Scheduling Jobs in XS Advanced. . . . .	1168
	Maintaining Users in XS Advanced. . . . .	1190
	Maintaining Database Instances in XS Advanced. . . . .	1197
	Maintaining the Service Broker in XS Advanced. . . . .	1200
	Maintaining the SAP HANA Deployment Infrastructure (HDI) Configuration . . . . .	1202
<b>9</b>	<b>SAP HANA Data Provisioning. . . . .</b>	<b>1218</b>
9.1	SAP HANA Smart Data Access. . . . .	1218
	Setting Up Database Drivers. . . . .	1219
	Creating and Configuring Remote Data Sources. . . . .	1232
	Using the Generic Adapter Framework. . . . .	1239
	Setting Up Single Sign-On (SSO) with Kerberos. . . . .	1243
	Creating Virtual Tables. . . . .	1247
	Creating Statistics on Virtual Tables. . . . .	1251
	Monitor Remote Connections and Remote Statements. . . . .	1253
	Remote Connection Pooling. . . . .	1256

	Smart Data Access System Parameters. . . . .	1257
9.2	SAP HANA Hadoop Integration. . . . .	1258
	Hadoop Integration Platform Support. . . . .	1259
	SAP HANA Spark Controller. . . . .	1259
	SAP HANA Ambari Integration. . . . .	1260
	Data Aging with Hadoop. . . . .	1261
	SAP HANA Vora. . . . .	1261
	Virtual Functions. . . . .	1261
	Using the Simba ODBC Driver to Connect to Hive. . . . .	1265
9.3	Data Replication and Transformation. . . . .	1268
	Replicating Data. . . . .	1269
	Transforming Data Using SAP HANA Web-based Development Workbench. . . . .	1280
	Transforming Data Using SAP HANA Application Function Modeler. . . . .	1291
	Node Reference. . . . .	1322
<b>10</b>	<b>SAP HANA HDBSQL (Command-Line Reference). . . . .</b>	<b>1346</b>
10.1	SAP HANA HDBSQL Options. . . . .	1346
10.2	Log On to a Database. . . . .	1352
10.3	Run Commands. . . . .	1353
10.4	Run Long Commands in Multiple-Line Mode. . . . .	1356
10.5	Edit Long Commands in an External File. . . . .	1357
10.6	Redirect Results to a File. . . . .	1358

# 1 Getting Started

The SAP HANA Administration Guide describes the main tasks and concepts necessary for the ongoing operation of SAP HANA.

## **i** Note

The *SAP HANA Administration Guide* does not cover administration tasks related to SAP HANA options and capabilities, such as SAP HANA dynamic tiering and SAP HANA smart data streaming. For more information about the administration of options and capabilities, see *SAP HANA Options and Capabilities* on SAP Help Portal. 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 \[page 1360\]](#).

The *SAP HANA Administration Guide* comprises the following main sections:

- **Overview of SAP HANA Architecture**

This section provides an overview of the basic architecture of SAP HANA systems:

- SAP HANA services
- Multiple-host systems
- Multitenant database containers

- **SAP HANA Administration Tools**

This section describes the main tools for the administration of SAP HANA: SAP HANA cockpit and SAP HANA studio.

- **System Administration**

This section describes all basic database administration and monitoring tasks including:

- Starting and stopping systems
- Creating, monitoring, and managing multitenant database containers
- System configuration (INI files)
- Monitoring (for example, performance, memory usage, disk usage, alert situations)
- Diagnosing and troubleshooting problems and error situations

- **Lifecycle Management**

This section covers the following aspects of lifecycle management:

- Platform lifecycle management
- Application lifecycle management
- Configuring SAP HANA applications with the process engine
- Deploying SAP HANA content

- **Security Administration**

This section describes administration tasks related to user provisioning, encryption, auditing, and in-database certificate management.

- **Availability and Scalability**

This section describes tasks related to the availability of your systems and the security of your data in failure situations. It is divided into the following main areas:

- High availability
- Backup and recovery

- 
- Scaling your systems
  - **Maintaining the Application Services Run-Time Environment**  
This section covers the administration tools used to manage the SAP HANA XS run-time environment for XS classic and XS advanced applications.
  - **SAP HANA Data Provisioning**  
This section covers tasks related to the following data acquisition features in SAP HANA:
    - SAP HANA Smart Data Access (SDA)
    - Hadoop Integration
    - Smart Data Integration (data replication and data transformation services)
  - **SAP HANA HDBSQL (Command Line Reference)**  
This section describes SAP HANA HDBSQL, a command line tool for executing commands on SAP HANA databases

For more information about the SAP HANA landscape, security, and installation, see SAP Help Portal.

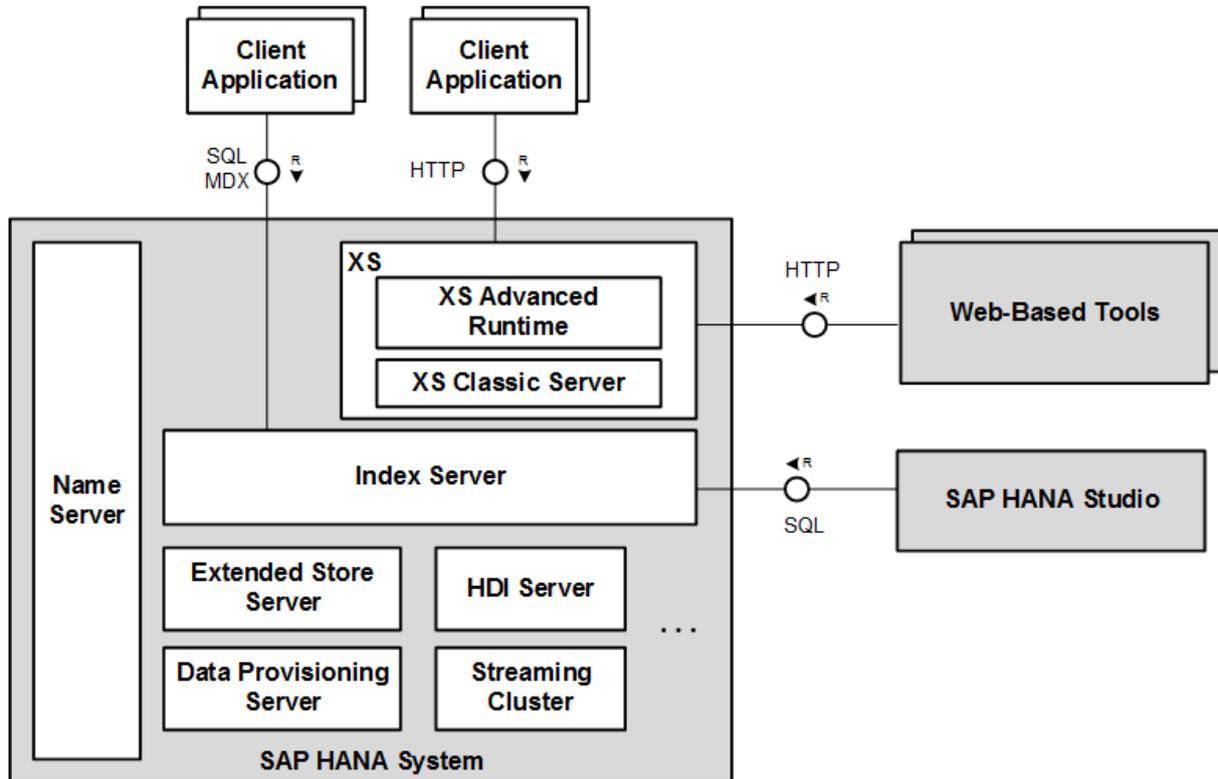
## Related Information

[SAP HANA Platform](#)

[SAP HANA Options and capabilities](#)

## 2 Overview of SAP HANA Architecture

An SAP HANA system comprises all the server components of an installation of SAP HANA. An SAP HANA system consists of several servers, the most important of which is the index server. The index server contains the actual data stores and the engines for processing the data.



Main Components of the SAP HANA System

The following is a brief overview of the most important server components of the SAP HANA system and the corresponding OS processes and services:

Server Component	OS Process	Service Name	Description
Index server	hdbindexserver	indexserver	The index server contains the actual data stores and the engines for processing the data.
Name server	hdbnameserver	nameserver	The name server owns the information about the topology of the SAP HANA system. In a distributed system with instances of the SAP HANA database on multiple hosts, the name server knows where the components are running and which data is located on which server.

Server Component	OS Process	Service Name	Description
XS classic server	hdbxsengine	xsengine	<p>SAP HANA Extended Application Services (SAP HANA XS) is the application server for native SAP HANA-based web applications. It is installed with the SAP HANA system 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>SAP HANA 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>
XS advanced runtime	<ul style="list-style-type: none"> <li>• hdbxscontroller</li> <li>• hdbxsxeagent</li> <li>• hdixsuaaserver</li> </ul>	<ul style="list-style-type: none"> <li>• xscontroller</li> <li>• xsexeagent</li> <li>• hdixsuaaserver</li> </ul>	<p>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.</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>
Extended store server	hdbesserver	esserver	<p>The extended store server is part of the SAP HANA dynamic tiering option of SAP HANA. It provides a high-performance disk-based column store for very big data up to the petabyte range.</p> <p>For more information about SAP HANA dynamic tiering, see <i>SAP HANA Options and Capabilities</i> on SAP Help Portal.</p>

Server Component	OS Process	Service Name	Description
Data provisioning server	hdbdpserver	dpserver	<p>The data provisioning server is part of the SAP HANA smart data integration option of SAP HANA. 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 about SAP HANA smart data integration, see <i>SAP HANA Options and Capabilities</i> on SAP Help Portal.</p>
Streaming cluster	hdbstreamingserver	streamingserver	<p>The streaming cluster is part of the SAP HANA smart data streaming option of SAP HANA. Smart data streaming extends SAP HANA with capabilities of SAP Event Stream Processor for consuming data streams and complex event processing.</p> <p>For more information about SAP HANA smart data streaming, see <i>SAP HANA Options and Capabilities</i> on SAP Help Portal.</p>
SAP HANA Deployment Infrastructure (HDI) server	hdbdiserver	diserver	HDI handles the deployment of design-time artifacts into SAP HANA.

### **i** Note

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 and Options \[page 1360\]](#).

In addition to the main servers mentioned above, the following auxiliary servers also run in the SAP HANA system:

Server Component	OS Process	Service Name	Description
Preprocessor server	hdbpreprocessor	preprocessor	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.
Compile server	hdbcompileserver	compileserver	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.
Script server	hdbscriptserver	scriptserver	The script server is used to execute application function libraries written in C++. The script server is optional and must be started manually. For more information, see SAP Note 1650957.
SAP Web Dispatcher	hdbwebdispatcher	webdispatcher	The Web Dispatcher processes inbound HTTP and HTTPS connections to XS services.

Server Component	OS Process	Service Name	Description
SAP start service	sapstartsrv	sapstartsrv	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.

## SAP HANA Extended Services Advanced Model

If the runtime platform of SAP HANA XS advanced is installed in your system, the following additional services run in the system:

Server Component	OS Process	Service Name	Description
SAP HANA XS Controller	hdbxscontroller	xscontroller	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.
SAP HANA XS Execution Agent	hdbxsxeagent	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.

For more information about SAP HANA XS advanced, see the *SAP HANA Development Guide (For SAP HANA XS Advanced Model)*.

### ➔ Recommendation

SAP recommends that customers and partners who want to develop new application use SAP HANA XS advanced model. If you want to migrate existing XS classic applications to run in the new XS advanced runtime 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.

## Related Information

[SAP Note 1650957](#)

[SAP HANA Options and Capabilities](#)

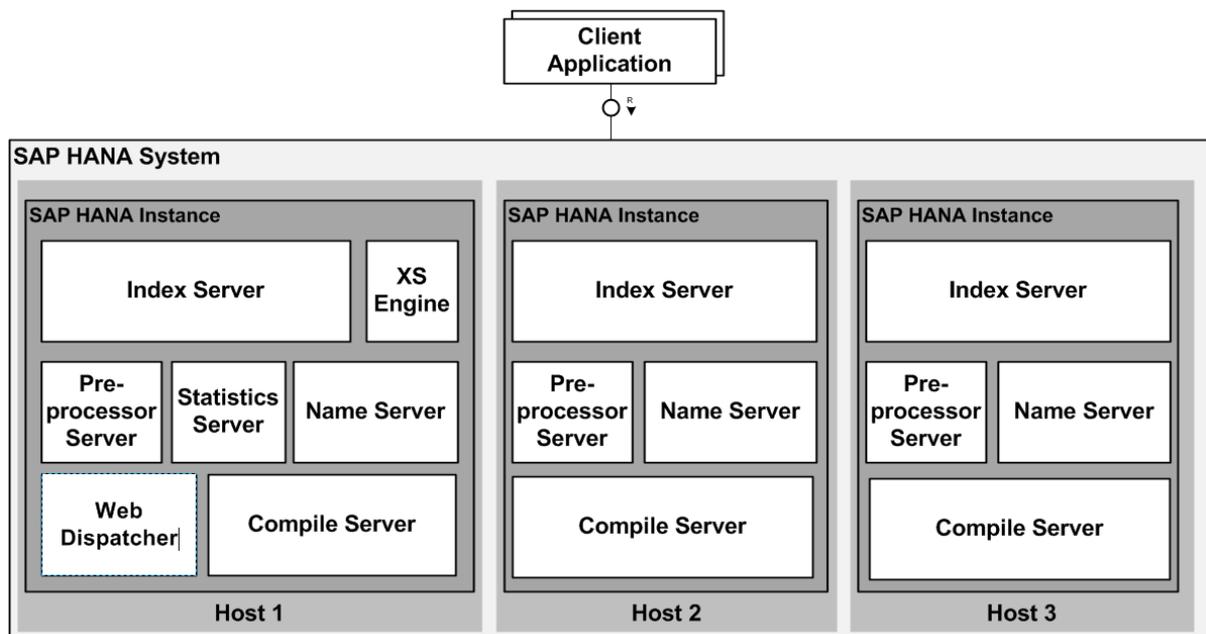
## 2.1 Multiple-Host Systems

SAP HANA supports the distribution of its server components 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. 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. In a multiple-host system, each index server is usually assigned its own host for maximum performance. It is possible to assign different tables to different hosts (partitioning the database), to split a single table between hosts (partitioning of tables), and to replicate tables to multiple hosts. For more information about hosts (for example, roles, fail-over configuration, and storage options), see *Multiple-Host System Concepts*.

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, shut down, or backup the system as a whole. The different server components of the system share the same metadata and requests from client applications can be transparently dispatched to different servers in the system.

The following figure shows a distributed system with three hosts, which each run a name server, index server, and so on.



Multiple-Host System

For more information about hosts (for example, roles, fail-over configuration, and storage options), see *Multiple-Host System Concepts* in the *SAP HANA Server Installation and Update Guide*.

### Related Information

[High Availability for SAP HANA \[page 774\]](#)

---

[Scaling SAP HANA \[page 998\]](#)

[Multiple-Host System Concepts \[page 532\]](#)

## 2.2 Multitenant Database Containers

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

An SAP HANA system installed in multiple-container mode is capable of containing more than one multitenant database containers. Otherwise, it is a single-container system. Single-container systems can be converted to multiple-container systems.

A multiple-container system always has exactly one system database, used for central system administration, and any number of multitenant database containers (including zero), also called tenant databases. An SAP HANA system installed in multiple-container mode 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 studio, connect to specific databases.

All the databases in a multiple-container system 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.

### 2.2.1 Server Architecture of Multiple-Container Systems

An SAP HANA system consists of multiple servers: name server, index server, preprocessor server, XS server and so on. The databases in a multiple-container system run different combinations of these servers.

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. Unlike the name server in a single-container system, the name server of the system database in a multiple-container 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. The SAP HANA XS

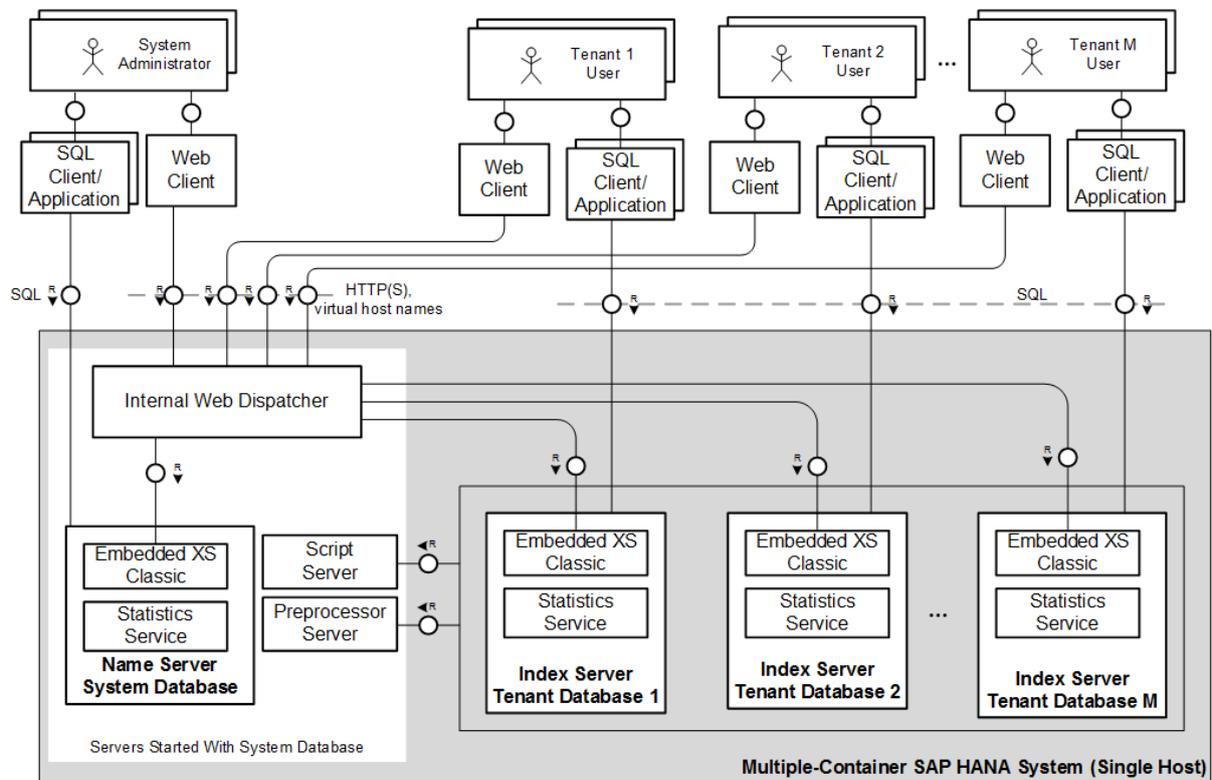
classic server runs embedded in the (master) index server of the tenant database by default, although it can be added as a separate service if necessary. Multiple-container systems are not yet supported by the SAP HANA XS advanced model.

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.

### Note

In addition to the system-internal Web Dispatcher, you can implement an external Web Dispatcher for load distribution. See *Using SAP Web Dispatcher for Load Balancing with Tenant Databases*.

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



Single-Host SAP HANA System with Multitenant Database Containers

## Related Information

[Using SAP Web Dispatcher for Load Balancing with Tenant Databases \[page 182\]](#)

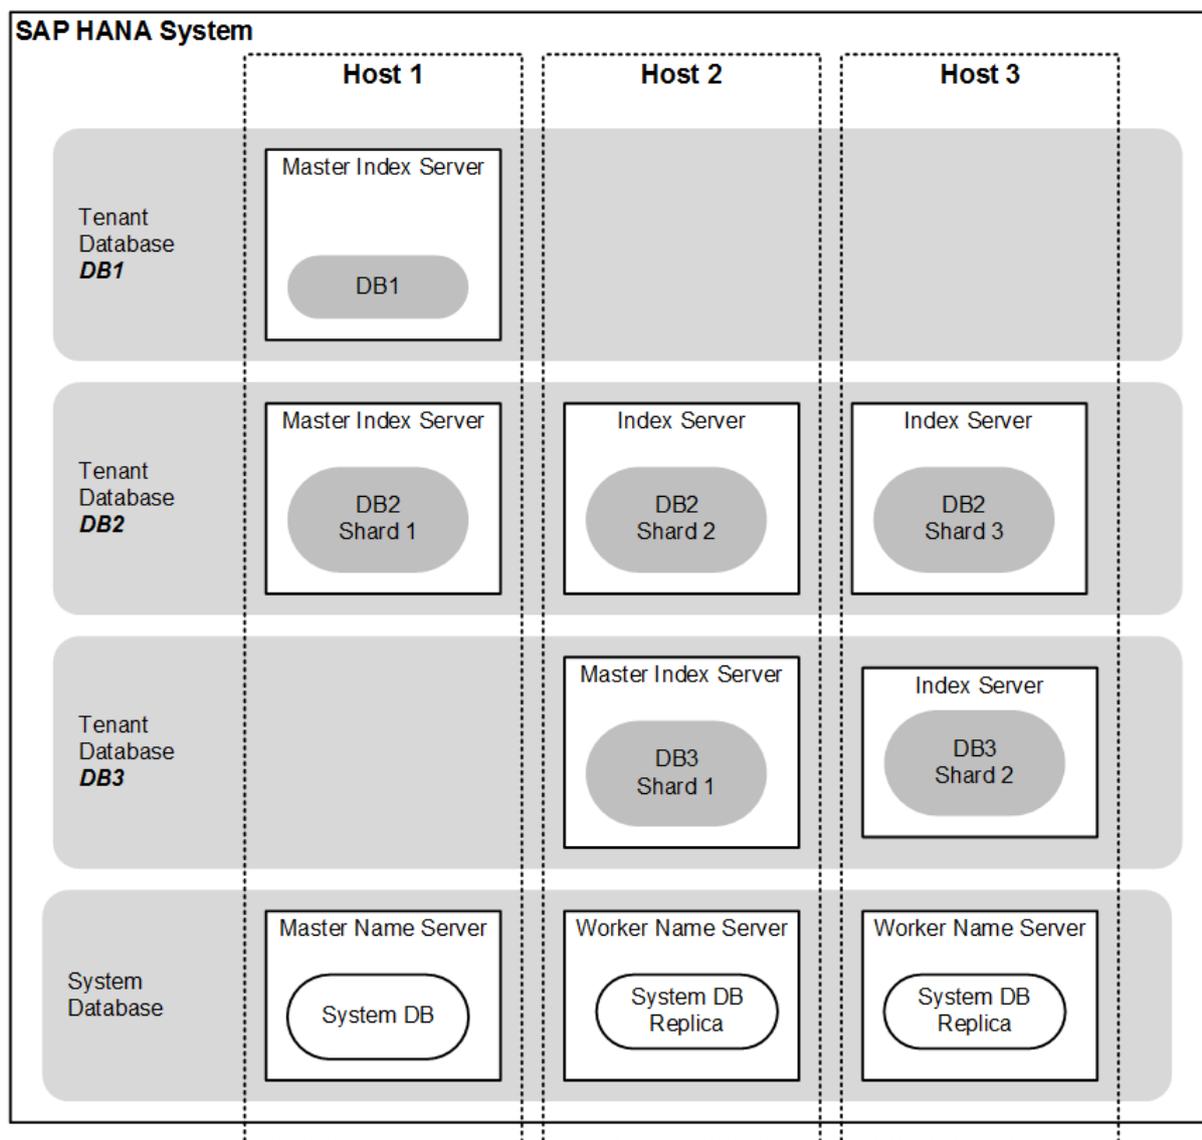
[Scale-Out Architecture of Multiple-Container Systems \[page 17\]](#)

## 2.2.2 Scale-Out Architecture of Multiple-Container Systems

It is possible to distribute tenant databases across several hosts (scale-out 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.

The following figure shows a multiple-container system with three tenant databases distributed across three hosts.



Multiple-Host System with Multitenant Database Containers

### Scale-Out Recommendations

When planning an SAP HANA MDC deployment, 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 \[page 998\]](#)

### 2.2.3 The System Database

The system database is created during either installation of a multiple-container system or conversion from a single-container system to a multiple-container 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 multiple-container system has exactly one system database. It is created during system installation or migration from a single-container system. It contains the data and users for system administration. System administration tools, such as the SAP HANA cockpit or the SAP HANA studio, 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 Multitenant Database Containers*.

---

## Things to Remember About the System 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 in a multiple-container system, you always have to create at least one tenant database.
- 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.

## Related Information

[Administration of Multitenant Database Containers \[page 19\]](#)

[Managing Resources in Multiple-Container Systems \[page 143\]](#)

## 2.2.4 Administration of Multitenant Database Containers

In SAP HANA systems that support multitenant database containers, there is a distinction between administration tasks performed at system level and those performed at database level.

### System Versus Database Administration

Unlike a single-container system in which system and database are perceived as a single unit and are therefore administered as one, multiple-container 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 (which is comparable to the administration of a single-container system), system-level administration of tenant databases requires the SAP HANA cockpit (for example, monitoring availability of tenant databases, creating and deleting tenant databases). The SAP HANA studio is required for system configuration tasks.

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

## Related Information

[Multitenant Database Containers \[page 15\]](#)

[The System Database \[page 18\]](#)

[Creating and Configuring Tenant Databases \[page 104\]](#)

[SAP HANA Administration Tools \[page 21\]](#)

[Monitoring and Managing Tenant Databases \[page 129\]](#)

### 3 SAP HANA Administration Tools

Several tools are available for the administration of SAP HANA.

Tool	Description
SAP HANA cockpit and SAP HANA cockpit for offline administration	<p>The SAP HANA cockpit is an SAP Fiori Launchpad site that provides you with a single point-of-access to a range of Web-based applications for the online administration of SAP HANA. You access the SAP HANA cockpit through a Web browser.</p> <p>To facilitate the execution of administration tasks as operating system user &lt;sid&gt;adm, a version of the SAP HANA cockpit that uses SAP Host Agent is available. The SAP HANA cockpit for offline administration allows you to perform administration tasks, such as starting the system or troubleshooting a system experiencing performance problems.</p>
SAP HANA studio	The SAP HANA Administration Console perspective of the SAP HANA studio supports general system administration and monitoring tasks.
SAP HANA lifecycle management tools	Specific tasks related to the configuration of system components after installation (such as adding and removing hosts, renaming systems) are performed using the SAP HANA lifecycle management tools.
SAP Solution Manager	If you are using SAP HANA in conjunction with other SAP business applications, it is possible to integrate with SAP Solution Manager.
SAP HANA HW Configuration Check Tool	This tool allows you to check the interoperability of SAP HANA with your existing enterprise storage in production environments.
SAP HANA XS Administration Tools	Both the SAP HANA XS classic model SAP HANA XS advanced model include a Web-based tool to enable you to maintain important parts of the application-development environment, for example, security and authentication methods.
SAP HANA Application Lifecycle Management	<p>This Web-based tool enables you to set up the transport of delivery units, start and monitor transports, and upload or download delivery unit archives.</p> <div style="background-color: #fff9c4; padding: 10px; margin-top: 10px;"> <p><b>i Note</b></p> <p>This tool is documented in the <i>SAP HANA Developer Guide (For SAP HANA Studio)</i>.</p> </div>
SAP HANA HDBSQL	SAP HANA HDBSQL is a command line tool for executing commands on SAP HANA databases.
SAP Landscape Virtualization Management software, enterprise edition	The enterprise edition of SAP Landscape Virtualization Management enables you to perform various operations on SAP HANA instances, such as start, stop, monitor, or system replication operations. For more information, see SAP Help Portal at <a href="http://help.sap.com/nwlvsm">http://help.sap.com/nwlvsm</a> > <a href="#">Application Help</a> > <a href="#">Managing System Landscapes</a> > <a href="#">Performing Operations on Instances</a> .

## Related Information

[SAP HANA Studio \[page 60\]](#)

[SAP HANA Cockpit \[page 22\]](#)

[Hardware Checks for Tailored Data Center Integration \[page 447\]](#)

[SAP HANA Platform Lifecycle Management \[page 500\]](#)

[Maintaining the Application Services Run-Time Environment \[page 1017\]](#)

[SAP Solution Manager for SAP HANA \[page 449\]](#)

[SAP HANA Application Lifecycle Management \[page 599\]](#)

[SAP Landscape Virtualization Management 2.1](#)

## 3.1 SAP HANA Cockpit

The SAP HANA cockpit is an SAP Fiori Launchpad site that provides you with a single point-of-access to a range of Web-based applications for the online administration of SAP HANA. You access the SAP HANA cockpit through a Web browser.

The SAP HANA cockpit is installed with SAP HANA as automated content. The applications to which it provides access are installed as separate delivery units, either as automated or non-automated content. In general, applications that provide core system and database administration features are available by default, for example, database monitoring, user management, and data backup. Other applications 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.

### Caution

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 \[page 1360\]](#).

The SAP HANA cockpit displays content as tiles arranged in groups. Tiles not only function as entry points to individual applications but also display selected app-specific data for immediate review. From these tiles, you can drill down into the relevant app for more detailed information and functions.

You can customize the default homepage of the SAP HANA cockpit by modifying the content of existing groups and creating your own groups. You can remove tiles that you don't need, as well as add tiles from any of the available tile catalogs.

The SAP HANA cockpit implements a role-based concept so that users only have access to those tiles for which they are authorized.

The SAP HANA cockpit is accessible at `https://<host>:43<instance>/sap/hana/admin/cockpit` (recommended) or `http://<host>:80<instance>/sap/hana/admin/cockpit`. For more information, see *Open SAP HANA Cockpit*.

## **i** Note

Full administration of SAP HANA is not possible with the SAP HANA cockpit. Certain tasks require the SAP HANA studio. For more information about which tasks are possible with the cockpit, see the documentation of the individual tile catalogs in *Related Information*.

## SAP HANA Cockpit for Offline Administration

To facilitate the execution of administration tasks as operating system user `<sid>adm`, a version of the SAP HANA cockpit that uses SAP Host Agent is available. The SAP HANA cockpit for offline administration allows you to perform administration tasks, such as starting the system or troubleshooting a system experiencing performance problems.

The SAP HANA cockpit for offline administration is accessible at `https://<host>:1129/lms1/hdbcockpit/<SID>/index.html` (recommended) or `http://<host>:1128/lms1/hdbcockpit/<SID>/index.html`. You can also navigate to the SAP HANA cockpit for offline administration from the standard SAP HANA cockpit.

For more information, see *SAP HANA Cockpit for Offline Administration*.

## Related Information

[Open SAP HANA Cockpit \[page 23\]](#)

[Tile Catalogs for Administration and Monitoring \[page 33\]](#)

[Tile Catalogs for Security Administration \[page 45\]](#)

[Tile Catalog for Platform Lifecycle Management \[page 52\]](#)

[SAP HANA Cockpit for Offline Administration \[page 56\]](#)

[SAP HANA Studio \[page 60\]](#)

## 3.1.1 Open SAP HANA Cockpit

You access the SAP HANA cockpit from a Web browser.

### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring` or `sap.hana.admin.roles::Administrator`.  
These privileges allow you to open the cockpit and access the tiles in the *SAP HANA Database Administration* catalog.

If you're opening the cockpit on the system database of a multiple-container, you also need the privileges granted by the role `sap.hana.admin.cockpit.roles::SysDBAdmin` so that you can access the tiles in the *SAP HANA System Administration* catalog.

You can grant roles using the *Assign Roles* app of the SAP HANA cockpit. For more information, see *Assign Roles to a User* in the *SAP HANA Administration Guide*.

### **i** Note

After database creation, you will have to log on for first time as the user SYSTEM. In this case, the aforementioned roles will be granted automatically. For more information, see *Roles Granted to Database User SYSTEM*.

- Your Web browser supports the SAPUI5 library `sap.m` (for example, Internet Explorer 9). For more information about SAPUI5 browser support, see SAP Note 1716423 and the Product Availability Matrix (PAM) for SAPUI5.
- If you're integrating the SAP HANA cockpit into a single sign-on (SSO) environment, you have specified and configured the methods for user authentication in the SAP HANA XS Administration Tool. Here, create a runtime configuration for the following applications:
  - `sap.hana.admin`
  - `sap.uis`For more information about how to do this, see *Maintaining the SAP HANA XS Classic Model Run Time* in the *SAP HANA Administration Guide*.
- For secure communication between SAP HANA and your browser, you have configured HTTPS. For more information, see *Maintaining HTTP Access to SAP HANA* in the *SAP HANA Administration Guide*.
- If you're opening the SAP HANA cockpit for either the system database or a tenant database in a multiple-container system, you have configured the internal SAP Web Dispatcher so that it can dispatch HTTP requests coming into the system to the correct database on the basis of DNS alias host names. Every tenant database needs an alias. For more information, see *Configure HTTP Access to Multitenant Database Containers* in the *SAP HANA Administration Guide*.

## Procedure

1. Enter the SAP HANA cockpit URL in your browser.

The URL depends on whether you are connecting to a single-container system or to a database in a multiple-container system.

A single-container system is accessed through the URL `https://<host_FQDN>:43<instance>/sap/hana/admin/cockpit` (recommended) or `http://<host_FQDN>:80<instance>/sap/hana/admin/cockpit`.

For more information about the URLs in multiple-container systems, see *Configure HTTP Access to Multitenant Database Containers* in the *SAP HANA Administration Guide*.

### **i** Note

FQDN = fully qualified domain name

2. If required, enter your database user name and password.

---

## Results

The SAP HANA cockpit opens. Groups for which you are authorized are displayed.

The name of the system or database appears in the shell bar of the launchpad. If the system has the system usage type `production`, this also indicated. For more information, see *Configure System Usage Type* in the *SAP HANA Administration Guide*.

## Next Steps

- Customize the homepage of the SAP HANA cockpit.
- If necessary, configure access to further tile catalogs and groups.

## Related Information

[Roles Granted to Database User SYSTEM \[page 56\]](#)

[Configure Access to Content in SAP HANA Cockpit \[page 31\]](#)

[Assign Roles to a User \[page 717\]](#)

[Configure HTTPS \(SSL\) for Client Application Access \[page 1067\]](#)

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

[SAP HANA XS Administration Tools \[page 1018\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Open SAP HANA Cockpit from SAP HANA Studio \[page 25\]](#)

[Monitoring in SAP HANA Cockpit \[page 290\]](#)

[Configure System Usage Type \[page 221\]](#)

[SAP Note 1716423 !\[\]\(912eb35f342458fc87c7c1d0cfd433ba\_img.jpg\)](#)

[Product Availability Matrix \(PAM\) for SAPUI5 !\[\]\(9f2eb39b5cb6ca001ddfe685f3184b1d\_img.jpg\)](#)

## 3.1.2 Open SAP HANA Cockpit from SAP HANA Studio

The SAP HANA cockpit is a Web application that you access from a Web browser. However, you can also launch the SAP HANA cockpit directly from the SAP HANA studio.

## Prerequisites

- You have added the system in the *Systems* view of the SAP HANA studio.
- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring` or `sap.hana.admin.roles::Administrator`.

These privileges allow you to open the cockpit and access the tiles in the [SAP HANA Database Administration](#) catalog.

If you're opening the cockpit on the system database of a multiple-container, you also need the privileges granted by the role `sap.hana.admin.cockpit.roles::SysDBAdmin` so that you can access the tiles in the [SAP HANA System Administration](#) catalog.

You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the [SAP HANA Administration Guide](#).

### **i** Note

After database creation, you will have to log on for first time as the user SYSTEM. In this case, the aforementioned roles will be granted automatically. For more information, see [Roles Granted to Database User SYSTEM](#).

- Your Web browser supports the SAPUI5 library `sap.m` (for example, Internet Explorer 9). For more information about SAPUI5 browser support, see SAP Note 1716423 and the Product Availability Matrix (PAM) for SAPUI5.
- For secure communication between SAP HANA and your browser, you have configured HTTPS. For more information, see [Maintaining HTTP Access to SAP HANA](#) in the [SAP HANA Administration Guide](#).
- If you're integrating the SAP HANA cockpit into a single sign-on (SSO) environment, you have specified and configured the methods for user authentication in the SAP HANA XS Administration Tool. Here, create a runtime configuration for the following applications:
  - `sap.hana.admin`
  - `sap.uis`For more information about how to do this, see [Maintaining the SAP HANA XS Classic Model Run Time](#) in the [SAP HANA Administration Guide](#).
- If you're opening the SAP HANA cockpit for either the system database or a tenant database in a multiple-container system, you have completed the following configuration steps:
  - Configure the internal SAP Web Dispatcher so that can dispatch HTTP requests coming into the system to the correct database on the basis of DNS alias hostnames. Every tenant database needs an alias. For more information, see [Configure HTTP Access to Multitenant Database Containers](#).
  - In the SAP HANA studio, specify the alias host name in the system properties of the tenant database ([▮ XS Properties](#) [▸ XS Host](#) [▮](#)). You don't have to change anything if you're connecting to the system database.

## Procedure

1. In the context menu of the [Systems](#) view, choose [▮ Configuration and Monitoring](#) [▸ Open SAP HANA Cockpit](#) [▮](#).  
SAP HANA cockpit opens in an external browser.
2. If required, enter your database user name and password.  
SAP HANA cockpit opens. The available content depends on your authorization.

---

## Next Steps

Customize the homepage of SAP HANA cockpit.

## Related Information

[Add an SAP HANA System \[page 70\]](#)

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

[SAP HANA XS Administration Tools \[page 1018\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Monitoring in SAP HANA Cockpit \[page 290\]](#)

[SAP Note 1716423](#)

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

## 3.1.3 Customize the Homepage of SAP HANA Cockpit

You can customize the default homepage of the SAP HANA cockpit by modifying existing groups and creating your own groups. It is also possible to configure the refresh interval of tiles that display certain monitoring data.

### 3.1.3.1 Change a Group

The homepage of the SAP HANA cockpit is delivered with default groups containing a default set of tiles in a default arrangement. You can modify these default groups, as well as other groups that you have created, by adding, removing, and rearranging tiles. You can also change the name of the group.

#### Procedure

1. Open the homepage in edit mode by clicking the pencil icon on the bottom right of the homepage.
2. To remove a tile from a group or to move to it another group, click the tile and choose *Remove* or *Move*.
3. To add a tile to a group:
  - a. Click the + tile.  
The tile catalog opens.
  - b. Find the tile you want.

#### ➔ Remember

You can only see tiles in the catalogs that you are authorized to use. If you can't see the catalog(s) that you need even though you have the correct role, the catalog may not be assigned to the role. For more information about how to do this, see *Configure Access to Content in SAP HANA Cockpit*.

- c. Choose the *Add to group* icon on the tile and select the group.
4. To rearrange the order of the tiles in a group, simply drag and drop them as you wish.
5. To change the title of a group, click the title and edit.
6. Return to the homepage by choosing the home icon on top left of the page and exit edit mode by clicking the pencil icon.

## Results

The group appears on your homepage containing only the required tiles.

## Related Information

[Configure Access to Content in SAP HANA Cockpit \[page 31\]](#)

### 3.1.3.2 Create a Group

You can create new groups to appear on the homepage of the SAP HANA cockpit. You can add tiles from any of the available tile catalogs to your new group(s).

## Procedure

1. Open the homepage in edit mode by clicking the pencil icon on the bottom right of the homepage.
2. Click *Add Group* wherever you want to add a group on the homepage.
3. Enter the name of the group.
4. Add tiles to the group:
  - a. Click the + tile.  
The tile catalog opens.
  - b. Find the tiles you want.

#### ➔ Remember

You can only see tiles in the catalogs that you are authorized to use. If you can't see the catalog(s) that you need even though you have the correct role, the catalog may not be assigned to the role. For more information about how to do this, see *Configure Access to Content in SAP HANA Cockpit*.

- c. Choose the *Add to group* icon on each tile and select the new group.
5. Return to the homepage by choosing the home icon on top left of the page and exit edit mode by clicking the pencil icon.

---

## Results

The new group is visible on the homepage.

## Related Information

[Configure Access to Content in SAP HANA Cockpit \[page 31\]](#)

### 3.1.3.3 Rename a Group

You can rename any group that appears on the homepage of the SAP HANA cockpit.

#### Procedure

1. Open the homepage in edit mode by clicking the pencil icon on the bottom right of the homepage.
2. Find the group you want to rename
3. Double-click the name and enter the new name.
4. Exit edit mode by clicking the pencil icon.

## Results

The group appears on the homepage with the new name.

### 3.1.3.4 Delete/Reset a Group

You can delete any group that you created on the homepage of the SAP HANA cockpit. You can't delete a default group but if you changed it in any way, you can reset it.

#### Procedure

1. Open the homepage in edit mode by clicking the pencil icon on the bottom right of the homepage.
2. Click *Delete* or *Reset*.

---

## Results

If you deleted the group, it no longer appears on your homepage. If you reset the group, the default tiles and title are restored.

### 3.1.3.5 Configure Automatic Tile Refresh

Tiles not only function as entry points to individual applications but also display selected app-specific data for immediate review, for example, current memory usage of the database. This data is refreshed automatically at default intervals, but if necessary, you can switch off automatic refresh or adjust the intervals.

## Context

To ensure that you have an up-to-date picture of the SAP HANA database, the information displayed on the tiles of the SAP HANA cockpit are automatically refreshed at default intervals. Intervals vary according to the type of information displayed. You can see which data is refreshed and how often in the homepage settings.

## Procedure

1. From the shell bar of the SAP HANA cockpit, choose **>> <username> > Settings >**.
2. Stop data on all tiles from being automatically refreshed by switching the *Automatic Tile Refresh* switch control to *No*.
3. Change the refresh interval a tile by selecting the required value.
4. Save and apply the new settings by choosing *OK*.

### 3.1.3.6 User Preferences

The *User Preferences* dialog provides read-only user connection information.

*User Preferences* appears as an entry in the options menu on the shell bar. The *User Preferences* dialog contains the following attributes.

#### **i** Note

All attributes are read only.

User Preference	Description
<user name>	Database user of the connected user
<i>Server</i>	Host and port of the connected system
<i>Language</i>	Browser interface language  Regardless of the language indicated here, the SAP HANA cockpit is available only in English. Some menus and dialogs from the underlying SAP Fiori framework may appear in the browser interface language.
<i>Theme</i>	Visual theme of the application  The SAP HANA cockpit uses the SAP standard theme <code>SAP Blue Crystal</code> .

### 3.1.4 Configure Access to Content in SAP HANA Cockpit

Before authorized end users can access content in the SAP HANA cockpit, tile catalogs and groups must be assigned to the relevant roles. For catalogs and groups delivered as default content, this is done automatically for the roles delivered with that content. However, for own-developed content or if you're not using standard roles, it needs to be done explicitly.

#### Prerequisites

- You have the following authorization:
  - Privileges granted by the role `sap.hana.uis.db::SITE_DESIGNER`  
This role allows you to edit Fiori launchpad application sites and catalogs.
  - EXECUTE privilege on the procedure `GRANT_ACTIVATED_ROLE (_SYS_REPO)`  
This privilege allows you to see all roles in the repository.

Roles and privileges can be granted in the *User* editor of the SAP HANA studio or the *Security* tool of the SAP HANA Web-based Workbench. For more information see *Create and Authorize a User* in the *SAP HANA Administration Guide*.

- The *Configure Role-Based Cockpit Access* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA User Management* tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

#### Context

In the SAP HANA cockpit, content is made available through tiles. Tile catalogs and groups contain tiles related to a particular functional area. For example, the *SAP HANA Database Administration* tile catalog and group contain tiles for database administration and monitoring.

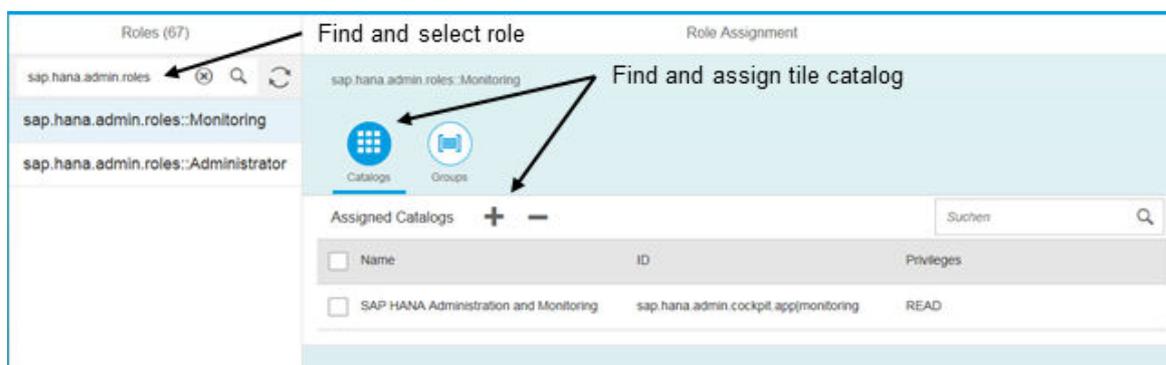
To be able to see the information and access the apps provided by tiles, end users must be assigned specific roles. Not only that, but the related tile catalogs and groups must also be assigned to these roles. Only then can authorized users see the tiles in the SAP HANA cockpit.

For tile catalogs and groups delivered as default content, this assignment is done automatically for the roles delivered with that content. However, for own-developed content or if you're not using the roles delivered with default content, it needs to be done explicitly. So, if authorized users cannot see the tile catalogs that they are authorized to use, do this assignment as described below.

## Procedure

1. Open the *Role Assignment* app by clicking the *Configure Role-Based Cockpit Access* tile on the homepage of the SAP HANA cockpit  
The *Role Assignment* app opens in a new browser window. The side panel displays a list of the roles available in the database, and the content area displays content assignments of the currently selected role, organized by tabs.  
The *Role Assignment* app opens in a new browser window. The side panel displays a list of the roles available in the database, and the content area displays content assignments of the currently selected role, organized by tabs.
2. Find and select the role with missing assignments.
3. In the content area, assign the related catalog(s) and group(s) to the role.

Role Assignment App



4. Close the *Role Assignment* app.

## Results

Users with the role can access the tiles of the assigned tile catalog and the group in the SAP HANA cockpit.

## Related Information

[Assign Roles to a User \[page 717\]](#)

[SAP Note 1716423](#)

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

## 3.1.5 Tile Catalogs for Administration and Monitoring

The tile catalog defines the set of all tiles available in the SAP HANA cockpit. Within the main tile catalog, tiles are grouped into catalogs according to functional area. A number of catalogs contain tiles for basic administration and monitoring.

### Tile Catalog: SAP HANA System Administration [page 33]

The catalog *SAP HANA System Administration* contains tiles that provide information and functions that allow you to monitor and manage tenant databases in a multiple-container system.

### Tile Catalog: SAP HANA Database Administration [page 35]

The catalog *SAP HANA Database Administration* contains tiles that provide information and functions related to database administration and monitoring.

### Tile Catalog: SAP HANA Backup [page 39]

The catalog *SAP HANA Backup* contains the *Data Backup* tile. This tile displays the status and time of the last completed data backup, and allows you to create and schedule database backups.

### Tile Catalog: SAP HANA System Replication [page 40]

The catalog *SAP HANA System Replication* contains the *System Replication* tile. This tile indicates whether or not the system is part of a system replication configuration. If it is, the tile provides you with the *System Replication* app where you can monitor the status of replication between the primary system and the secondary system(s).

### Tile Catalog: Smart Data Access Administration [page 42]

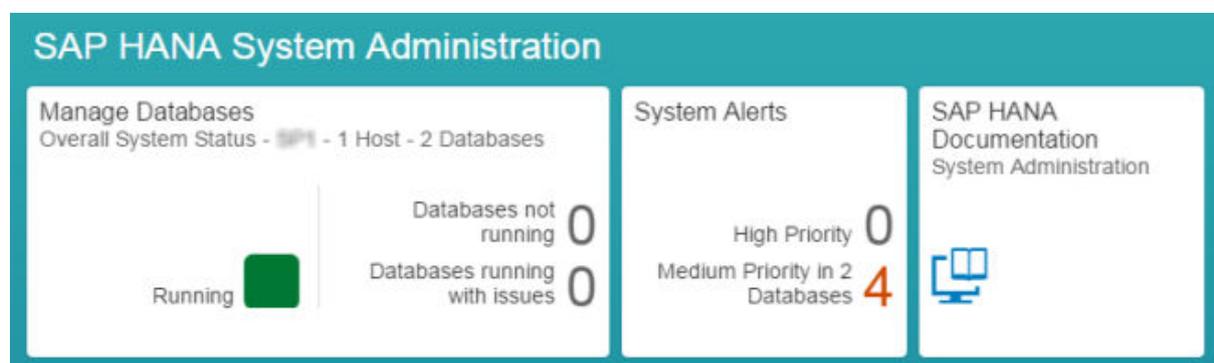
The catalog *Smart Data Access Administration* contains the *Remote Connections Monitor* and *Remote Statements Monitor* tiles. These tiles provide monitoring information related to Smart Data Access (remote sources and virtual tables).

### Tile Catalog: SAP HANA Performance Management [page 44]

The *SAP HANA Performance Management* catalog contains tiles that provide information and functions, which allow you to capture, replay, and analyze the workload from an SAP HANA system.

### 3.1.5.1 Tile Catalog: SAP HANA System Administration

The catalog *SAP HANA System Administration* contains tiles that provide information and functions that allow you to monitor and manage tenant databases in a multiple-container system.



Tile Catalog SAP HANA System Administration

## **i** Note

This tile catalog is only available in the system database. It is not available if you access the SAP HANA cockpit from a single-container system or from a tenant database in a multiple-container system.

The tiles available with this catalog, as well as the roles assigned to it, are described in the following sections.

## Tiles

The information available on each tile as well as drill-down options are described in the following table:

Tile	Description
<a href="#">Manage Databases</a>	<p>Indicates overall system health</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"><li>• Running All databases are running.</li><li>• Running with issues All databases are running, but some have high-priority alerts.</li><li>• Not running One or more databases are not running.</li></ul> <p>This tile opens the <a href="#">Manage Databases</a> app where you can monitor the status and resource usage of individual databases (with drill-down), as well as perform other administration tasks such as stopping and starting tenant databases and creating new tenant databases.</p>
<a href="#">System Alerts</a>	<p>Indicates the number of high and medium alerts currently raised in tenant databases and opens the <a href="#">Alerts</a> app where you can view and analyze alert details</p>
<a href="#">SAP HANA Documentation – System Administration</a>	<p>Opens the SAP HANA documentation that describes those system-level administration tasks that you can perform using the SAP HANA cockpit</p>

## ➔ Remember

If you're accessing the SAP HANA cockpit from the system database of a multiple-container system, the tiles in the [SAP HANA System Administration](#) tile catalog provide information and functions for monitoring and managing all tenant databases in the system. If you want to administer the system database directly, use the tiles in the [SAP HANA Database Administration](#) catalog.

## Roles

The privileges in the following roles are required to access the tiles in this catalog:

Role	Description
<code>sap.hana.admin.cockpit.sysdb.roles::SysDBAdmin</code>	Allows users in the system database to monitor and manage tenant databases in a multiple-container system

### ➔ Recommendation

Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. 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.

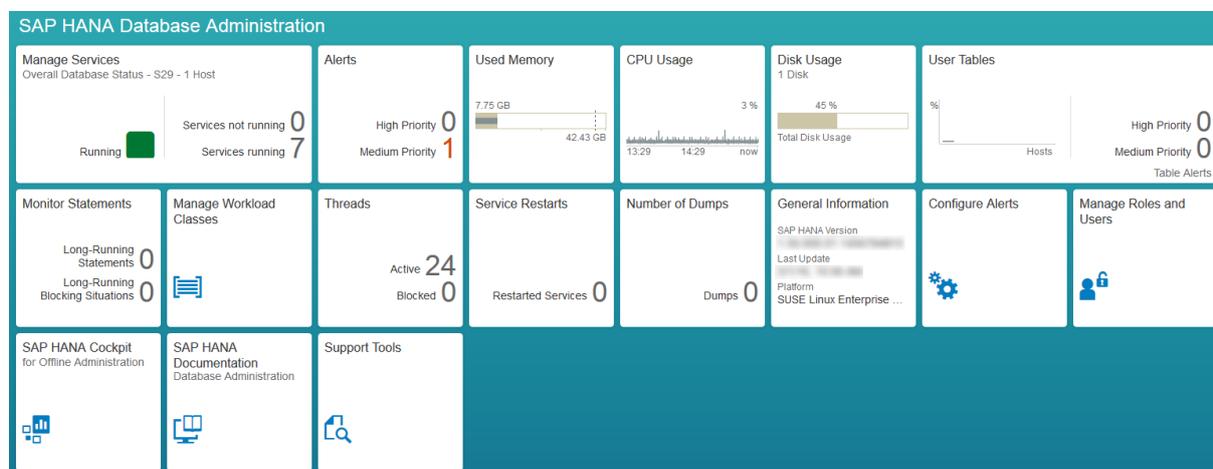
## Related Information

[Managing Multitenant Database Containers \[page 104\]](#)

### 3.1.5.2 Tile Catalog: SAP HANA Database Administration

The catalog *SAP HANA Database Administration* contains tiles that provide information and functions related to database administration and monitoring.

Tile Catalog SAP HANA Database Administration



The tiles available with this catalog, as well as the roles assigned to it, are described in the following sections.

## Tiles

The information available on each tile as well as drill-down options are described in the following table:

Tile	Description
<a href="#">Manage Services</a>	<p>Indicates overall database health</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> <li>• Running All services are running.</li> <li>• Running with issues All services are running, but there are high-priority alerts.</li> <li>• Not running One or more services are not started.</li> </ul> <p>The number of services running and not running are indicated. If the database is distributed across multiple hosts, this includes all services on worker hosts.</p> <p>This tile opens the <a href="#">Manage Services</a> app where you can monitor the status and resource usage of individual database services, as well as perform other administration tasks such as stopping and removing services.</p>
<a href="#">Alerts</a>	<p>Indicates the number of high and medium alerts currently raised in the database and opens the <a href="#">Alerts</a> app where you can view and analyze alert details</p>
<a href="#">Used Memory</a>	<p>Indicates the total amount of memory currently used by the SAP HANA database in relation to the allocation limit</p> <p>If the database is distributed across multiple hosts, the memory usage of all worker hosts is indicated. The host with the highest (most critical) memory usage is shown in more detail.</p> <p>The bars indicating memory usage change color (gray, orange, and red) based on the thresholds of the alert "Host physical memory usage" (ID 2).</p> <p>This tile opens the <a href="#">Performance Monitor</a> app where you can visualize and explore the usage history of key system resources (CPU, memory, and disk). When you access the <a href="#">Performance Monitor</a> app from the <a href="#">Used Memory</a> tile, memory-related KPIs are automatically selected.</p>
<a href="#">CPU Usage</a>	<p>Indicates the percentage of CPU used by the SAP HANA database compared with the operating system as a whole</p> <p>If the database is distributed across multiple hosts, the CPU usage of all worker hosts is indicated. The host with the highest (most critical) CPU usage is shown in more detail.</p> <p>The percentage value indicating CPU usage changes color (gray, orange, and red) based on the thresholds of the alert "Host CPU usage" (ID 5).</p> <p>This tile opens the <a href="#">Performance Monitor</a> app where you can visualize and explore the usage history of key system resources (CPU, memory, and disk). When you access the <a href="#">Performance Monitor</a> app from the <a href="#">CPU Usage</a> tile, CPU-related KPIs are automatically selected.</p>

Tile	Description
<a href="#">Disk Usage</a>	<p>Indicates the <b>total</b> usage of all disks, that is including space used by non-SAP HANA data</p> <p>The disk with the highest (most critical) disk usage is shown in more detail.</p> <p>The bars indicating disk usage change color (gray, orange, and red) based on the thresholds of the alert "Disk usage" (ID 2).</p> <p>This tile opens the <a href="#">Performance Monitor</a> app where you can visualize and explore the usage history of key system resources (CPU, memory, and disk). When you access the <a href="#">Performance Monitor</a> app from the <a href="#">Disk Usage</a> tile, disk-related KPIs are automatically selected.</p>
<a href="#">User Tables</a>	<p>Indicates the number of hosts monitored and the name of the host with the highest memory usage.</p> <p>If alerts exist, the tile displays the total for medium and high priority alerts.</p> <p>This tile opens the <a href="#">User Tables</a> app where you can visualize tables by size, explore the usage history of tables, and move tables to warm storage.</p>
<a href="#">Monitor Statements</a>	<p>Indicates the number of long-running statements and blocking situations as determined by the corresponding alerts (39 and 49)</p> <p>If statement memory tracking is enabled, this is also indicated.</p> <p>This tile opens the <a href="#">Monitor Statements</a> app where you can analyze the most critical statements currently running in the database and if necessary, enable statement memory tracking.</p>
<a href="#">Manage Workload Classes</a>	<p>This tile opens the <a href="#">Manage Workload Classes</a> app.</p>
<a href="#">Threads</a>	<p>Indicates the number of currently active and blocked threads</p> <p>This tile opens the <a href="#">Threads</a> app where you can analyze active threads currently running in the database.</p>
<a href="#">Service Restarts</a>	<p>Indicates the number of services that have been manually or automatically restarted</p> <p>If restarts have been detected, this tile opens the <a href="#">Alerts</a> app where you can view the related alerts.</p>
<a href="#">Number of Dumps</a>	<p>Indicates the number of dump files in the database's trace directory</p> <p>This tile opens the <a href="#">Trace</a> tool of the SAP HANA Web-based Development Workbench.</p> <div data-bbox="603 1659 1396 1870" style="background-color: #fff9c4; padding: 10px;"> <p><b>i Note</b></p> <p>The <a href="#">Trace</a> tool of the SAP HANA Web-based Development Workbench opens in a new window and requires additional roles, either <code>sap.hana.xs.ide.roles::TraceViewer</code> or the parent role <code>sap.hana.xs.ide.roles::Developer</code>.</p> </div>
<a href="#">General Information</a>	<p>Provides information about system version, time of last upgrade, and platform</p> <p>This tile opens the <a href="#">General Information</a> app.</p>

Tile	Description
<a href="#">Configure Alerts</a>	Indicates if any alerts are failing to run and opens the <a href="#">Alert Configuration</a> app where you can configure alert schedules and thresholds and set up e-mail notification
<a href="#">Manage Roles and Users</a>	<p>Opens the <a href="#">Security</a> tool of the SAP HANA Web-based Development Workbench where you can provision users</p> <div style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;"> <p><b>i Note</b></p> <p>The <a href="#">Security</a> tool of the SAP HANA Web-based Development Workbench opens in a new window and requires additional roles, either <code>sap.hana.xs.ide.roles::SecurityAdmin</code> or the parent role <code>sap.hana.xs.ide.roles::Developer</code>.</p> </div>
<a href="#">SAP HANA Cockpit - for Offline Administration</a>	Opens the SAP HANA cockpit for offline administration where you can stop and (re)start the system, as well as troubleshoot and diagnose problems, even when the system is stopped or cannot be reached by SQL due to performance problems
<a href="#">SAP HANA Documentation – Database Administration</a>	Opens the SAP HANA documentation that describes those database administration tasks that you can perform using the SAP HANA cockpit
<a href="#">Support Tools</a>	This tile opens the <a href="#">Support Tools</a> app.

## Roles

The privileges in the following roles are required to access the tiles in this catalog:

Role	Description
<code>sap.hana.admin.roles::Monitoring</code>	<p>Allows users to open the SAP HANA cockpit with read-only access to monitoring data</p> <p>This role also allows users to see tiles in the <a href="#">SAP HANA Platform Lifecycle Management</a> and <a href="#">Smart Data Access Administration</a> tile catalogs.</p>
<code>sap.hana.admin.roles::Administrator</code>	<p>Allows users to open the SAP HANA cockpit with read-only access to monitoring data, as well as to perform database administration tasks supported by the cockpit (configure alerts, stop/start services, reset memory statistics, cancel sessions)</p> <p>This role also allows users to see tiles in the <a href="#">SAP HANA Platform Lifecycle Management</a> tile catalog.</p>

### ➔ Recommendation

Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. 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.

## Related Information

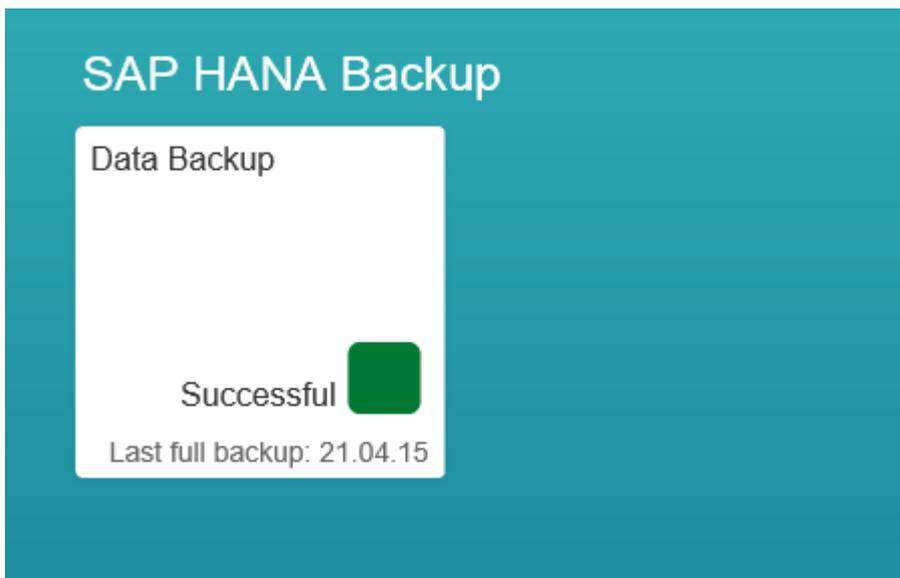
[Monitoring in SAP HANA Cockpit \[page 290\]](#)

[Tile Catalog for Platform Lifecycle Management \[page 52\]](#)

[Tile Catalog: Smart Data Access Administration \[page 42\]](#)

### 3.1.5.3 Tile Catalog: SAP HANA Backup

The catalog *SAP HANA Backup* contains the *Data Backup* tile. This tile displays the status and time of the last completed data backup, and allows you to create and schedule database backups.



SAP HANA Backup

## Roles

The privileges in the following roles are required to work with the backup functionality:

- Create backups and monitor the backup status.  
`sap.hana.backup.roles::Operator`  
`sap.hana.backup.roles::Administrator`
- Schedule backups.  
`sap.hana.backup.roles::Scheduler`  
Scheduling backups requires the XS Job Scheduler to be active.  
More information: *Scheduling Jobs in XS Advanced in Related Information*

### ➔ Recommendation

Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. 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.

## Related Information

[Create Data Backups and Delta Backups \(SAP HANA Cockpit\) \[page 926\]](#)

[Display Information About Backups \(SAP HANA Cockpit\) \[page 928\]](#)

[Schedule Data Backups \(SAP HANA Cockpit\) \[page 933\]](#)

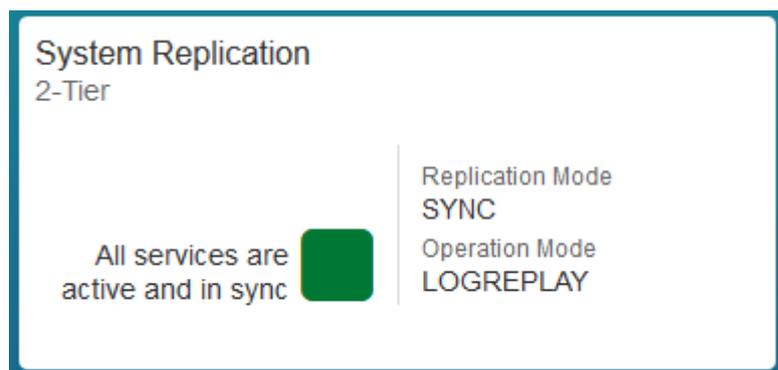
[SAP HANA Backup \[page 880\]](#)

[Backup Catalog \[page 902\]](#)

[Scheduling Jobs in XS Advanced \[page 1168\]](#)

## 3.1.5.4 Tile Catalog: SAP HANA System Replication

The catalog *SAP HANA System Replication* contains the *System Replication* tile. This tile indicates whether or not the system is part of a system replication configuration. If it is, the tile provides you with the *System Replication* app where you can monitor the status of replication between the primary system and the secondary system(s).



SAP HANA System Replication

The tiles available with this catalog, as well as the roles assigned to it, are described in the following sections.

## Tiles

The information available on the tile as well as drill-down options are described in the following table:

Tile	Description
<p><i>System Replication</i></p>	<p>Indicates status of system replication</p> <p>Information displayed includes:</p> <ul style="list-style-type: none"> <li>• Type of landscape (2 tier or 3 tier)</li> <li>• Replication mode (SYNC, SYNCMEM or ASYNC) – in a 3 tier landscape the replication modes of both sections are shown</li> <li>• Operation mode (LOGREPLAY or DELTA DATA SHIPPING)</li> <li>• The following states can be shown:               <ul style="list-style-type: none"> <li>○ Not configured (meaning system replication is not configured)</li> <li>○ Active and in sync (green square)</li> <li>○ All services are active but not yet in sync yet (yellow triangle)</li> <li>○ Errors in Replication (red circle)</li> </ul> </li> </ul> <p>This tile opens the System Replication app where you see an overview of the system replication: the header shows the number of active and standby hosts (for multi host systems with more than two hosts) / host name (for systems with less than 3 nodes), average log buffer write wait time, which depending on the replication mode in use, shows the time taken to ship the log buffers to the secondary.</p> <ul style="list-style-type: none"> <li>• SYNC/SYNCMEM: round trip time to send the log buffers and receive an acknowledgment</li> <li>• ASYNC: ASYNC: start time is when the log buffers are created, end time is when they are sent out to the network.</li> </ul> <p>All changes to data are captured in the redo log, which SAP HANA persists in form of log buffers of 4 KB to 1 MB size in the log volumes. In SAP HANA system replication every write transaction requires that the redo log buffers are not only written locally to persistent storage but are also shipped to the secondary site. This log buffer write wait time KPI represents the time taken to ship log buffers to the secondary site over the last 24 hours.</p>

## Roles

The privileges in the following roles are required to access the tiles in this catalog:

Role	Description
<code>sap.hana.admin.cockpit.sysrep.roles::SysRepAdmin</code>	Allows users read-only access to monitor system replication status

### ➔ Recommendation

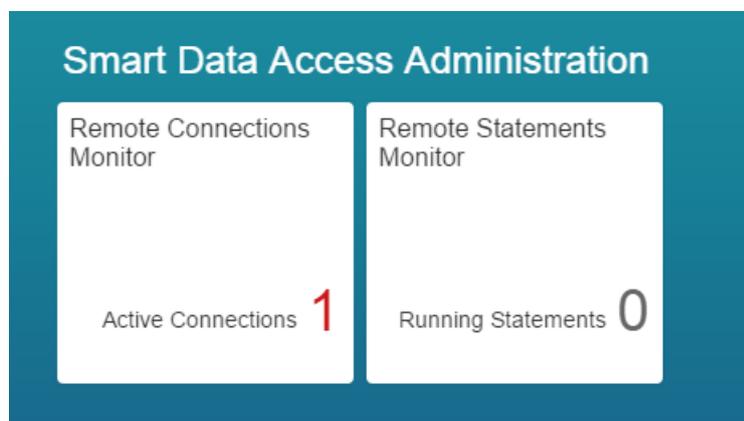
Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. Furthermore, if repository package privileges are granted by a role, we recommend that these privileges be restricted to your organization's packages rather than the complete 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.

## Related Information

[System Replication Details \[page 844\]](#)

### 3.1.5.5 Tile Catalog: Smart Data Access Administration

The catalog *Smart Data Access Administration* contains the *Remote Connections Monitor* and *Remote Statements Monitor* tiles. These tiles provide monitoring information related to Smart Data Access (remote sources and virtual tables).



Smart Data Access Administration

The tiles available with this catalog, as well as the roles assigned to it, are described in the following sections.

## Tiles

The information available on each tile as well as drill-down options are described in the following table:

Tile	Description
Remote Connections Monitor	Indicates the number of active remote connections  This tile opens the <i>Remote Connections Monitor</i> app, where you can analyze remote connections in the database.
Remote Statements Monitor	Indicates the number of running remote statements  This tile opens the <i>Remote Statements Monitor</i> app, where you can analyze remote statements in the database.

## Roles

The privileges in the following roles are required to access the tiles in this catalog:

Role	Description
<code>sap.hana.admin.roles::Monitoring</code>	Allows users to open the SAP HANA cockpit with read-only access to monitoring data.  This role also allows them to see the tiles in the <i>Smart Data Access Administration</i> tile catalog.

### ➔ Recommendation

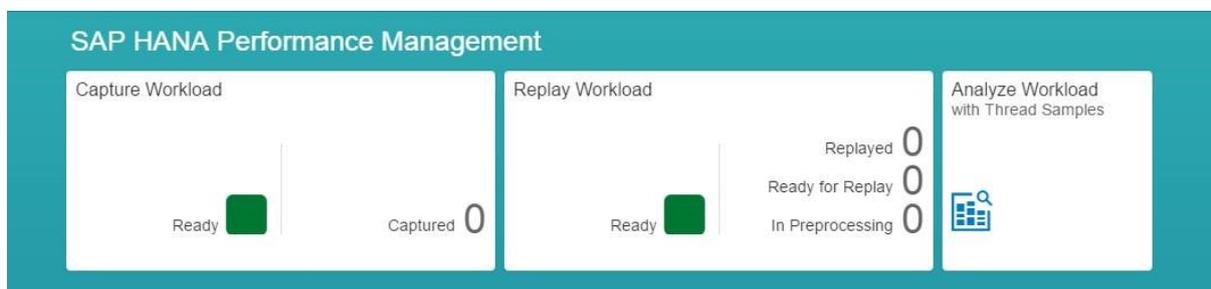
Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. Furthermore, if repository package privileges are granted by a role, we recommend that these privileges be restricted to your organization's packages rather than the complete 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.

## Related Information

[Monitor Remote Connections and Remote Statements \[page 1253\]](#)

### 3.1.5.6 Tile Catalog: SAP HANA Performance Management

The *SAP HANA Performance Management* catalog contains tiles that provide information and functions, which allow you to capture, replay, and analyze the workload from an SAP HANA system.



SAP HANA Performance Management

The tiles available with this catalog, as well as the roles assigned to it, are described in the following sections.

#### **i** Note

For *Capture Workload* and *Replay Workload* this catalog is only available if the delivery unit `HANA_WORKLOAD_REPLAY 1.0` has been deployed. For more information, see *Deploy a Delivery Unit Archive (\*.tgz)*.

## Tiles

The information available on the tile as well as drill-down options are described in the following table:

Tile	Description
<i>Capture Workload</i>	<p>Indicates the number of captured workloads</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> <li>• Ready</li> <li>• Capturing</li> </ul> <p>The tile opens the <i>Capture Workload</i> app, where you can capture and monitor workloads.</p>
<i>Replay Workload</i>	<p>Indicates the number of replayed workloads, workloads ready for replay, as well as workloads in preprocessing</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> <li>• Ready</li> <li>• Replaying</li> </ul> <p>The tile opens the <i>Replay Workload</i> app, where you can pre-process workloads, replay preprocessed workloads and monitor during workload replay.</p>

Tile	Description
<a href="#">Analyze Workload</a>	<p>Indicates overall system health</p> <p>The tile opens the <a href="#">Analyze Workload</a> app, where you can identify the root cause of performance issues.</p>

## Roles

The privileges in the following roles are required to access the tiles in this catalog:

Role	Description
<code>sap.hana.replay.roles::Capture</code>	Allows users to open the SAP HANA cockpit with access to the tiles in the <a href="#">SAP HANA Performance Management</a> tile catalog. This role also allows users to capture workloads..
<code>sap.hana.replay.roles::Replay</code>	Allows users to open the SAP HANA cockpit with access to the tiles in the <a href="#">SAP HANA Performance Management</a> tile catalog. This role also allows users to replay workloads..
<code>sap.hana.workloadanalyzer.roles::Operator</code>	Allows users to open the SAP HANA cockpit with access to the tiles in the <a href="#">SAP HANA Performance Management</a> tile catalog. This role also grants users read-only access to analyze workloads.

### ➔ Recommendation

Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. 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.

## Related Information

[Deploy a Delivery Unit Archive \(\\*.tgz\) \[page 632\]](#)

## 3.1.6 Tile Catalogs for Security Administration

The tile catalog defines the set of all tiles available in the SAP HANA cockpit. Within the main tile catalog, tiles are grouped into catalogs according to functional area. Two catalogs contain tiles for security administration tasks in SAP HANA.

[Tile Catalog: SAP HANA Security Overview \[page 46\]](#)

The *SAP HANA Security Overview* tile catalog contains information related to critical security settings.

[Tile Catalog: SAP HANA User Management \[page 49\]](#)

The *SAP HANA User Management* tile catalog contains information and functions related to user management.

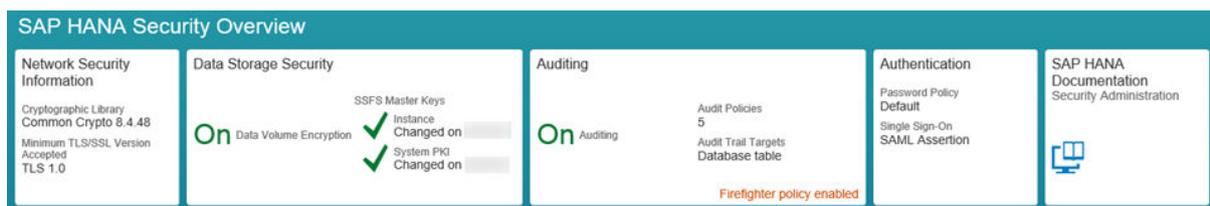
[Tile Catalog: SAP HANA Certificate Management \[page 50\]](#)

The catalog *SAP HANA Certificate Management* contains tiles that provide information and functions related to the management of X.509 client certificates stored in the database. These certificates are used in the SAP HANA database for certificate-based user authentication and secure client-server communication using the Secure Sockets Layer (SSL) protocol.

### 3.1.6.1 Tile Catalog: SAP HANA Security Overview

The *SAP HANA Security Overview* tile catalog contains information related to critical security settings.

Tile Catalog SAP HANA Security Overview



## Tiles

The information available on each tile as well as drill-down options are described in the following table:

Tile	Description
<a href="#">Network Security Information</a>	<p>Indicates the cryptographic library in use in the system and the minimum accepted version of the transport layer security/secure sockets layer (TLS/SSL) protocol</p> <p>This tile opens the <i>Network Security Information</i> app where you can see more detailed information about network configuration.</p>

Title	Description
<a href="#">Auditing</a>	<p>Indicates whether or not auditing is enabled in the system, the number of audit policies, the configured audit trail target, as well as any auditing-related alerts</p> <p>If a firefighter policy is active in the system (that is, a policy that audits all the actions of a particular user), this is also indicated.</p> <p>This tile opens the <a href="#">Auditing</a> app 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>
<a href="#">Data Storage Security</a>	<p>Indicates whether or not data volumes are encrypted, as well as when the master keys of the secure stores in the file system (SSFS) were changed</p> <p>This tile opens the <a href="#">Data Volume Encryption</a> app where you can see more information about the encryption status of individual data volume, enable or disable encryption, and change the root encryption key.</p> <div style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;"> <p><b>⚠ Caution</b></p> <p>Do not enable data volume encryption in an existing system without having first read the section <i>Enable Data Volume Encryption in an Existing SAP HANA System</i>.</p> </div>
<a href="#">Authentication</a>	<p>Indicates the status of the password policy (default or customized), as well as the user authentication mechanisms configured for single sign-on in the database</p> <p>This tile opens the <a href="#">Password Policy and Blacklist</a> app where you can see and edit the password policy and blacklist.</p>
<a href="#">SAP HANA Documentation – Security Administration</a>	Provides access to the SAP HANA documentation that describes those security administration tasks that you can perform using the SAP HANA cockpit

## Roles

The privileges in the following roles are required to access the tiles in this catalog:

Role	Description
<code>sap.hana.security.cockpit.roles::DisplaySecurityDashboard</code>	Allows users to see information about critical security settings (auditing, data storage, and network communication)
<code>sap.hana.security.cockpit.roles::MaintainDataVolumeEncryption</code>	Allows users to see information about critical security settings, as well as to enable and disable data volume encryption and change the root key used for data volume encryption

Role	Description
<code>sap.hana.security.cockpit.roles::MaintainPasswordPolicy</code>	Allows users to see information about critical security settings, as well as to edit the password policy and password blacklist
<code>sap.hana.security.cockpit.roles::MaintainAuditPolicy</code>	Allows users to see information about critical security settings, as well as to create and edit audit policies and to make global auditing settings such as enabling and disabling auditing and configuring audit trail targets

### ➔ Recommendation

Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. 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.

## Related Information

[View Status of Security Settings \[page 634\]](#)

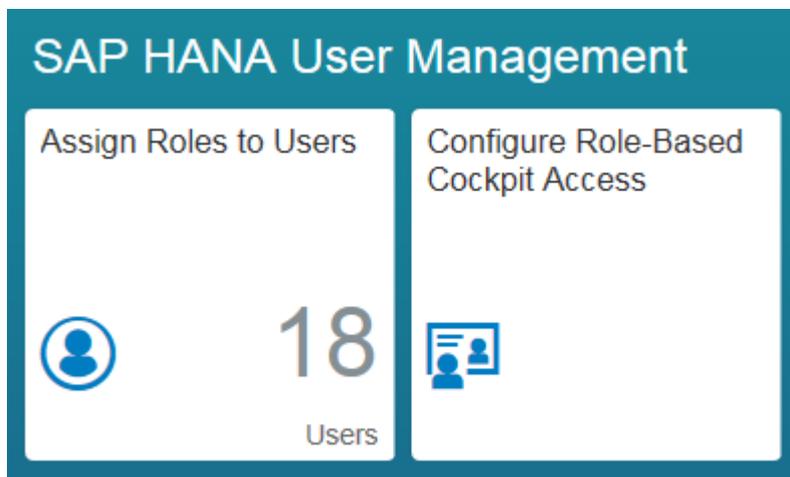
[Auditing Details \[page 724\]](#)

[Network Security Details \[page 635\]](#)

[Enable Data Volume Encryption in an Existing SAP HANA System \[page 750\]](#)

## 3.1.6.2 Tile Catalog: SAP HANA User Management

The *SAP HANA User Management* tile catalog contains information and functions related to user management.



Tile Catalog SAP HANA User Management

The tiles available with this catalog, as well as the roles assigned to it, are described in the following sections.

### Tiles

The information available on each tile as well as drill-down options are described in the following table:

Tile	Description
<i>Assign Roles to Users</i>	Indicates the number users in the database and provides you with access to the <i>Assign Roles</i> app where you can assign catalog and repository roles to users.
<i>Configure Role-Based Cockpit Access</i>	<p>Provides access to the <i>Role Assignment</i> app where you can tile catalog(s) and group(s) to roles</p> <p>This allows users with the roles to access the catalogs and groups.</p> <p><b>i Note</b></p> <p>This tile is not visible in the default group on the launchpad. You must add it from the catalog. For more information, see <i>Change a Group</i>.</p>

## Roles

The privileges in the following roles are required to access the tiles in this catalog:

Role	Description
<code>sap.hana.security.cockpit.roles::DisplayAssignedRoles</code>	Allows users to see which roles are granted to users
<code>sap.hana.security.cockpit.roles::EditAssignedRoles</code>	Allows users to grant roles to users

### ➔ Recommendation

Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. 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.

## Related Information

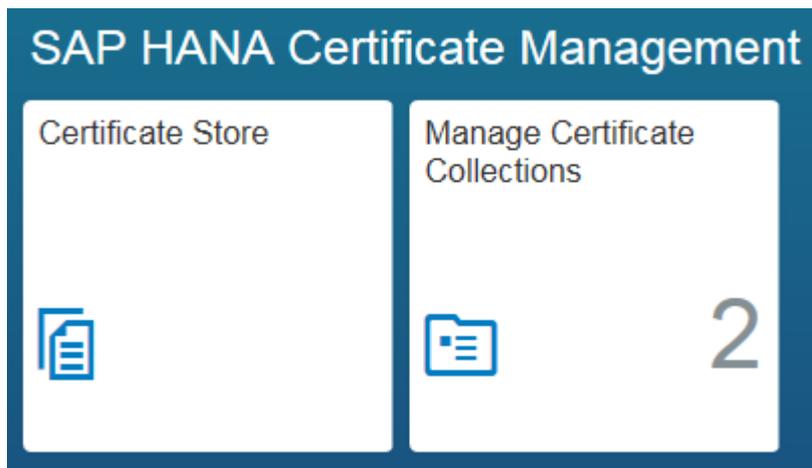
[Assign Roles to a User \[page 717\]](#)

[Configure Access to Content in SAP HANA Cockpit \[page 31\]](#)

[Change a Group \[page 27\]](#)

### 3.1.6.3 Tile Catalog: SAP HANA Certificate Management

The catalog *SAP HANA Certificate Management* contains tiles that provide information and functions related to the management of X.509 client certificates stored in the database. These certificates are used in the SAP HANA database for certificate-based user authentication and secure client-server communication using the Secure Sockets Layer (SSL) protocol.



Tile Catalog SAP HANA Certificate Management

The tiles available with this catalog, as well as the roles assigned to it, are described in the following sections.

## Tiles

The information available on each tile as well as drill-down options are described in the following table:

Tile	Description
<a href="#">Certificate Store</a>	Provides access to the certificate store, an in-database repository for X.509 client certificates
<a href="#">Configure Certificate Collections</a>	Indicates the number of collections and the number of certificates that are due to expire (if any), and provides access to the <a href="#">Certificate Collections</a> app

## Roles

The privileges in the following roles are required to access the tiles in this catalog:

Role	Description
<code>sap.hana.security.cockpit.roles::DisplayCertificateStore</code>	Allows users read-only access to certificates and certificate collections stored in the database
<code>sap.hana.security.cockpit.roles::MaintainCertificates</code>	Allows users to import trusted certificates into the certificate store
<code>sap.hana.security.cockpit.roles::MaintainCertificateCollections</code>	Allows users to create collections, as well as add trusted certificates and server certificates to collections
<code>sap.hana.security.cockpit.roles::EditCertificateStore</code>	Allows user to set the purpose of a collection in conjunction with either system privilege USER ADMIN or SSL admin and object privilege REFERENCES on the collection

### ➔ Recommendation

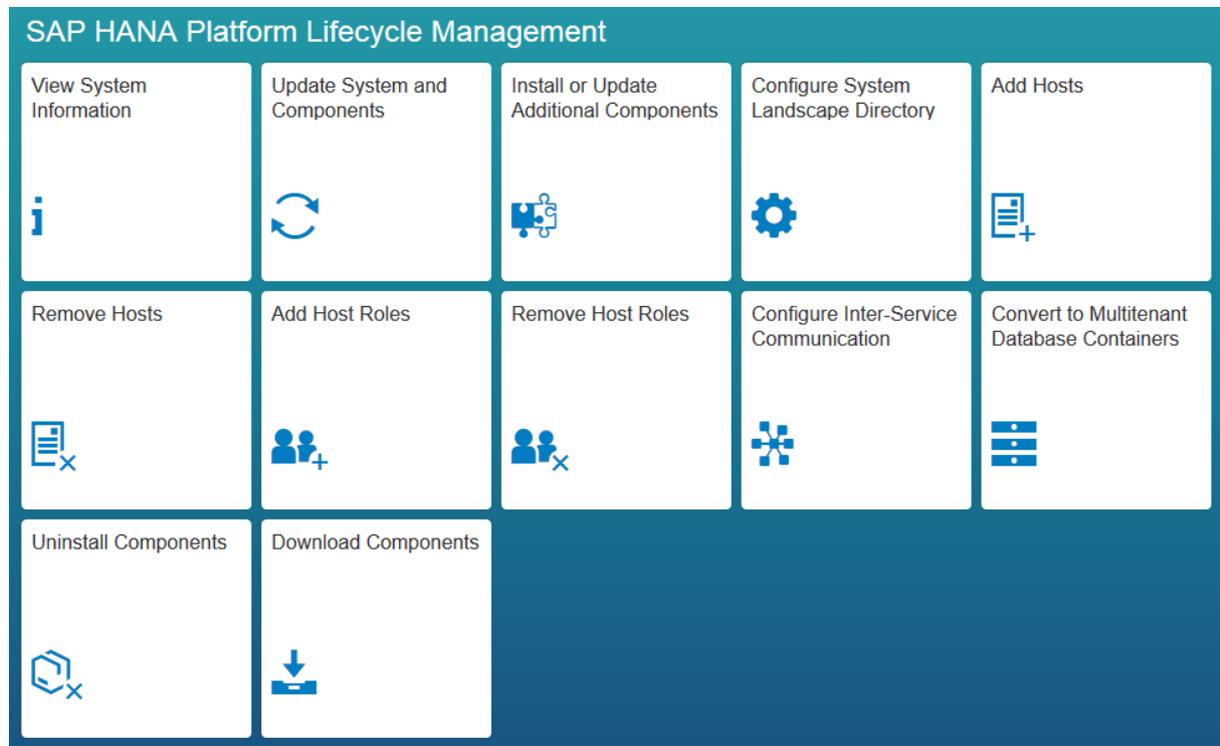
Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. Furthermore, if repository package privileges are granted by a role, we recommend that these privileges be restricted to your organization's packages rather than the complete 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.

## Related Information

[Managing Client Certificates in the SAP HANA Database \[page 758\]](#)

## 3.1.7 Tile Catalog for Platform Lifecycle Management

The tile catalog *SAP HANA Platform Lifecycle Management* contains tiles that provide information and functions related to platform lifecycle management.



Tile Catalog SAP HANA Platform Lifecycle Management

The tiles available with this catalog, as well as the roles assigned to it, are described in the following sections.

## Tiles

The information available on each tile as well as drill-down options are described in the following table:

Tile	Description
<a href="#">View System information</a>	Provides information about system, hosts, and installed components.
<a href="#">Update System and Components</a>	Allows users to update SAP HANA database server components.
<a href="#">Install or Update Additional Components</a>	Allows users to install additional 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).
<a href="#">Configure System Landscape Directory Registration</a>	Allows users to configure a connection to the System Landscape Directory (SLD).
<a href="#">Configure Inter-Service Communication</a>	Allows users to configure inter-service communication to prevent unauthorized access to the SAP HANA system via the internal communication channels in multiple-host systems.
<a href="#">Add Hosts</a>	Allows users to add one or more hosts to the system.
<a href="#">Remove Hosts</a>	Allows users to remove one or more hosts from the system.
<a href="#">Add Host Roles</a>	Allows users to add one or more host roles to hosts.
<a href="#">Remove Host Roles</a>	Allows users to remove one or more host roles from hosts.
<a href="#">Uninstall Components</a>	Allows users to uninstall one or more components.
<a href="#">Convert to Multitenant Database Containers</a>	Allows users to convert an SAP HANA system to support multitenant database containers.
<a href="#">Download Components</a>	Allows users to download components from the SAP Support Portal.

## Roles

The privileges in the following roles are required to access the tiles in this catalog:

Role	Description
<code>sap.hana.admin.roles::Monitoring</code>	<p>Allows users to open the SAP HANA cockpit with read-only access to monitoring data</p> <p>This role also allows users to see tiles in the <a href="#">SAP HANA Platform Lifecycle Management</a> and <a href="#">Smart Data Access Administration</a> tile catalogs.</p>

Role	Description
<code>sap.hana.admin.roles::Administrator</code>	<p>Allows users to open the SAP HANA cockpit with read-only access to monitoring data, as well as to perform database administration tasks supported by the cockpit (configure alerts, stop/start services, reset memory statistics, cancel sessions)</p> <p>This role also allows users to see tiles in the <a href="#">SAP HANA Platform Lifecycle Management</a> tile catalog.</p>

### ➔ Recommendation

Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. 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.

## Related Information

[SAP HANA Lifecycle Management \[page 499\]](#)

## 3.1.8 Roles for Tile Catalogs

To access the various tile catalogs of the SAP HANA cockpit, users must be assigned specific roles.

The following table lists the roles required to access individual tile catalogs out of the box.

### ➔ Recommendation

We recommend that you use these roles as templates for creating your own roles. You then need to reconfigure access to the tile catalog. See *Configure Access to Content in SAP HANA Cockpit*

To Access the Tile Catalog...	You Need the Privileges in the Role...
<p><a href="#">SAP HANA System Administration</a></p> <p><b>i Note</b></p> <p>Both the tile catalog <a href="#">SAP HANA System Administration</a> and the associated role are only available in the system database of a multiple-container system.</p>	<p><code>sap.hana.admin.cockpit.sysdb.roles::SysDBAdmin</code></p>

To Access the Tile Catalog...	You Need the Privileges in the Role...
<i>SAP HANA Database Administration</i>	<ul style="list-style-type: none"> <li>• <code>sap.hana.admin.roles::Monitoring</code>, OR</li> <li>• <code>sap.hana.admin.roles::Administrator</code></li> </ul>
<i>SAP HANA Platform Lifecycle Management</i>	<ul style="list-style-type: none"> <li>• <code>sap.hana.admin.roles::Monitoring</code>, OR</li> <li>• <code>sap.hana.admin.roles::Administrator</code></li> </ul>
<i>SAP HANA Backup</i>	<ul style="list-style-type: none"> <li>• <code>sap.hana.backup.roles::Operator</code>, OR</li> <li>• <code>sap.hana.admin.roles::Administrator</code></li> </ul>
<i>SAP HANA System Replication</i>	<code>sap.hana.admin.cockpit.sysrep.roles::SysRepAdmin</code>
<i>SAP HANA Security Overview</i>	<code>sap.hana.security.cockpit.roles::DisplaySecurityDashboard</code>
<i>SAP HANA User Management</i>	<ul style="list-style-type: none"> <li>• <code>sap.hana.security.cockpit.roles::DisplayAssignedRoles</code>, OR</li> <li>• <code>sap.hana.security.cockpit.roles::EditAssignedRoles</code></li> </ul>
<i>SAP HANA Certificate Management</i>	<ul style="list-style-type: none"> <li>• <code>sap.hana.security.cockpit.roles::DisplayCertificateStore</code>, OR</li> <li>• <code>sap.hana.security.cockpit.roles::MaintainCertificates</code>, OR</li> <li>• <code>sap.hana.security.cockpit.roles::MaintainCertificateCollections</code>, OR</li> <li>• <code>sap.hana.security.cockpit.roles::EditCertificateStore</code></li> </ul>
<i>SAP HANA Performance Management</i>	<ul style="list-style-type: none"> <li>• <code>sap.hana.replay.roles::Capture</code>, OR</li> <li>• <code>sap.hana.replay.roles::Replay</code>, OR</li> <li>• <code>sap.hana.workloadanalyzer.roles::Operator</code></li> </ul>

## Related Information

[Tile Catalog: SAP HANA System Administration \[page 33\]](#)

[Tile Catalog for Platform Lifecycle Management \[page 52\]](#)

[Tile Catalog: SAP HANA Database Administration \[page 35\]](#)

[Tile Catalog: SAP HANA Backup \[page 39\]](#)

[Tile Catalog: SAP HANA System Replication \[page 40\]](#)

[Tile Catalog: SAP HANA User Management \[page 49\]](#)

[Tile Catalog: SAP HANA Certificate Management \[page 50\]](#)

[Tile Catalog: SAP HANA Performance Management \[page 44\]](#)

[Configure Access to Content in SAP HANA Cockpit \[page 31\]](#)

## 3.1.9 Roles Granted to Database User SYSTEM

To ensure that SAP HANA cockpit can be used immediately after database creation, the database user SYSTEM is automatically granted several roles the first time the cockpit is opened with this user.

### Caution

Do not use the SYSTEM user for day-to-day activities. Instead, use this user to create dedicated database users for administrative tasks and to assign privileges to these users. It is recommended that you then deactivate the SYSTEM user. For more information see *Deactivate the SYSTEM User* in this guide.

Role	Description
<code>sap.hana.admin.roles::Administrator</code>	Allows users to open the SAP HANA cockpit with read-only access to monitoring data, as well as to perform database administration tasks supported by the cockpit (configure alerts, stop/start services, reset memory statistics, cancel sessions)  This role also allows users to see tiles in the <a href="#">SAP HANA Platform Lifecycle Management</a> tile catalog.
<code>sap.hana.xs.ide.roles::TraceViewer</code>	Allows users to open the <a href="#">Trace</a> tool of the SAP HANA Web-based Development Workbench
<code>sap.hana.ide.roles::SecurityAdmin</code>	Allows users to open the <a href="#">Security</a> tool of the SAP HANA Web-based Development Workbench
<code>sap.hana.admin.cockpit.sysdb.roles::SysDBAdmin</code>	Allows users in the system database to monitor and manage tenant databases in a multiple-container system

### Note

`sap.hana.admin.cockpit.roles::SysDBAdmin` is granted only if you are logging on to the system database of a multiple-container system.

## Related Information

[Deactivate the SYSTEM User \[page 640\]](#)

## 3.1.10 SAP HANA Cockpit for Offline Administration

The SAP HANA cockpit for offline administration is a version of the SAP HANA cockpit that communicates with SAP HANA using SAP Host Agent. You can use it to perform administration tasks, such as starting the system or troubleshooting a system experiencing performance problems, as operating system user `<sid>adm`. In a system replication scenario, you can use it to perform a takeover.

---

Similarly to the standard SAP HANA cockpit, the SAP HANA cockpit for offline administration displays content as tiles that function as entry points to individual applications. However, there are no groups or tile catalogs, and it is not possible to modify the content or the homepage.

➔ **Tip**

For more information about the standard SAP HANA cockpit, see *SAP HANA Cockpit*.

## Communication with SAP Host Agent

The SAP Host Agent is a tool used to perform several lifecycle management tasks in SAP systems. In SAP HANA, it is installed by default on all hosts.

Together with the SAP HANA database lifecycle manager (HDBLCM), the SAP HANA cockpit for offline administration relies on the SAP Host Agent to execute tasks as the system administrator user `<sid>adm`.

Both the SAP HANA database lifecycle manager and the SAP HANA cockpit for offline administration can communicate with the SAP Host Agent via HTTPS (port 1129) or HTTP (port 1128). The SAP HANA database lifecycle manager uses HTTPS by default and automatically handles the certificate configuration required for secure communication. HTTPS access to the SAP HANA cockpit for offline administration is therefore enabled without any further configuration necessary. For more information, see *Secure Sockets Layer (SSL) Certificate Handling*.

➔ **Recommendation**

We recommend accessing the SAP HANA cockpit for offline administration via HTTPS. Passwords are transferred in plain text via HTTP. Since it is possible to open the SAP HANA cockpit for offline administration from the standard SAP HANA cockpit, we also recommend using HTTPS for the standard SAP HANA cockpit.

For more information about the SAP Host Agent in an SAP HANA installation, see *Using SAP Host Agent to Execute Platform LCM Tasks*.

## Related Information

[SAP HANA Cockpit \[page 22\]](#)

[Using SAP Host Agent to Execute Platform LCM Tasks \[page 522\]](#)

[Secure Sockets Layer \(SSL\) Certificate Handling \[page 523\]](#)

[Perform a Takeover \[page 848\]](#)

## 3.1.10.1 Open SAP HANA Cockpit for Offline Administration

You access the SAP HANA cockpit for offline administration from a Web browser or from the standard SAP HANA cockpit.

### Prerequisites

- You have the credentials of the operating system user (<sid>adm user) that was created when the system was installed.
- Communication port 1129 is open.  
Port 1129 is required for communication with the SAP Host Agent in a standalone browser via HTTPS.
- You have configured the standard SAP HANA cockpit for HTTPS access.  
This is required for HTTPS access to the SAP HANA cockpit for offline administration from the standard SAP HANA cockpit. For more information, see *Configure HTTPS (SSL) for Client Application Access* in the *SAP HANA Administration Guide*.
- Your Web browser supports the SAPUI5 library `sap.m` (for example, Internet Explorer 9).  
For more information about SAPUI5 browser support, see SAP Note 1716423 and the Product Availability Matrix (PAM) for SAPUI5.

### Procedure

1. Open the SAP HANA cockpit for offline administration.

You can do this the following ways:

Option	Description
<b>Directly</b>	<p>Enter the URL in your browser:</p> <pre>https://&lt;host&gt;:1129/lmsl/hdbcockpit/&lt;sid&gt;/index.html</pre> <div style="background-color: #fff9c4; padding: 10px;"><p><b>i Note</b></p><p>It's also possible to access the SAP HANA cockpit for offline administration via the URL <code>http://&lt;host&gt;:1128/lmsl/hdbcockpit/&lt;sid&gt;/index.html</code>). However, this is not recommended because passwords are transferred in plain text via HTTP.</p></div>
<b>From the SAP HANA cockpit</b>	<ol style="list-style-type: none"><li>1. Open the SAP HANA cockpit by entering the URL in your browser: <code>https://&lt;host&gt;:43&lt;instance&gt;/sap/hana/admin/cockpit</code> (recommended) or <code>http://&lt;host&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code></li><li>2. In the <i>SAP HANA Database Administration</i> group, click the tile <i>SAP HANA Cockpit for Offline Administration</i>.</li></ol>

Option	Description
	<p><b>i Note</b></p> <p>If you access the SAP HANA cockpit via HTTP, then the SAP HANA cockpit for offline administration is also accessed via HTTP. Therefore, we recommend configuring the SAP HANA cockpit for HTTPS access.</p>

2. Enter the <sid>adm user name and password.

## Results

The SAP HANA cockpit for offline administration opens. For more information about the tiles displayed, see *Tiles for Offline Administration and Diagnosis*.

## Related Information

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

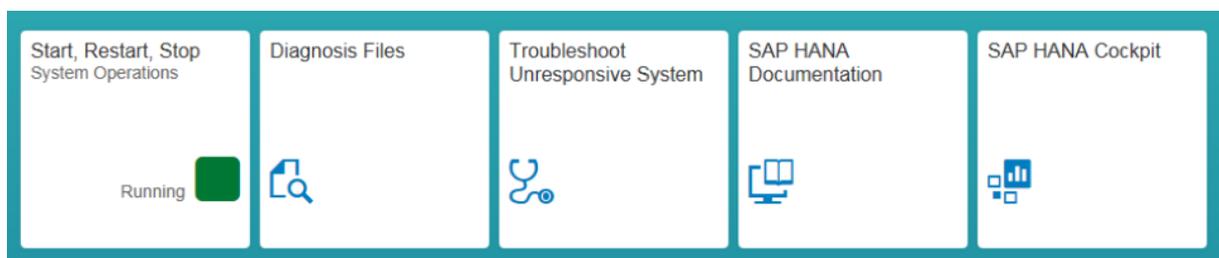
[SAP Note 1716423](#)

[Tiles for Offline Administration \[page 59\]](#)

[Configure HTTPS \(SSL\) for Client Application Access \[page 1067\]](#)

### 3.1.10.2 Tiles for Offline Administration

The homepage of the SAP HANA cockpit for offline administration contains several tiles. Some of the tiles are not available in a secondary system.



Tiles for Offline Administration

Tile	Description
<i>Start, Restart, Stop</i>	<p>Indicates the system status</p> <ul style="list-style-type: none"> <li>• Initializing</li> <li>• Running</li> <li>• Stopping</li> <li>• Stopped</li> </ul> <p>This tile opens the <i>System Operations</i> app where you can start, restart, and stop the system.</p>
<i>Diagnosis Files</i>	<p>Opens the <i>Diagnosis Files</i> app where you can access log, trace and other diagnosis files</p> <p>This app also allows you to trigger the collection of diagnosis information into a zip file, which you can then download and attach to a support message, for example.</p>
<i>Troubleshoot Unresponsive System</i>	<p>Triggers the collection of transactional information and then displays this information for troubleshooting performance issues</p> <p>If you are using SAP HANA system replication, this tile is not available in the SAP HANA cockpit of secondary systems.</p>
<i>SAP HANA Documentation - SAP HANA Offline Administration</i>	<p>Opens the SAP HANA documentation that describes those administration tasks that you can perform using the SAP HANA cockpit for offline administration</p>
<i>SAP HANA Cockpit</i>	<p>Opens the SAP HANA cockpit where you can access all applications for the online administration of SAP HANA</p> <p>If you are using SAP HANA system replication, this tile is not available in the SAP HANA cockpit of secondary systems.</p>

## Related Information

[Starting and Stopping Systems in SAP HANA Cockpit \[page 91\]](#)

[View Diagnosis Files in SAP HANA Cockpit \[page 463\]](#)

[Collect and Download Diagnosis Information in SAP HANA Cockpit \[page 486\]](#)

[Troubleshoot an Unresponsive System in SAP HANA Cockpit \[page 482\]](#)

## 3.2 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 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.

#### **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.

## 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.

### 3.2.1 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**

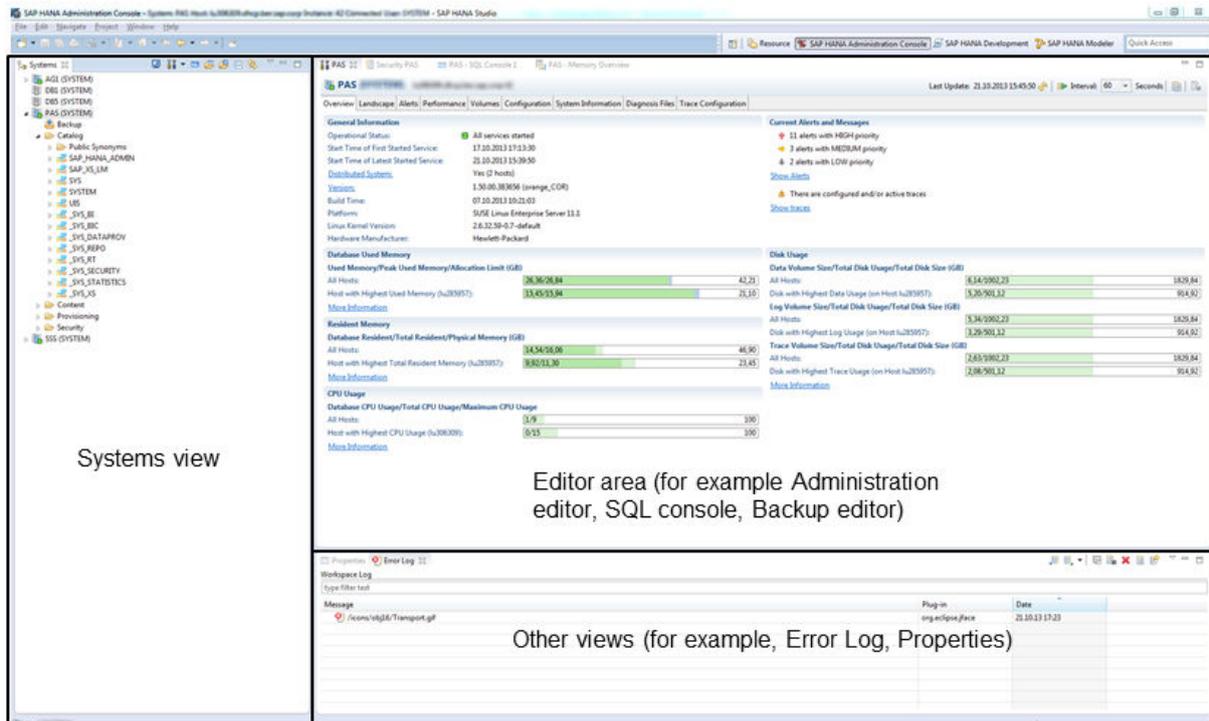
Administration Console or by choosing the  SAP HANA Administration Console button in the perspective switcher in the upper-right corner of the screen.

### 3.2.1.1 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"> <li>• The System Monitor</li> <li>• The Administration editor</li> <li>• The Backup editor</li> <li>• The Security editor</li> <li>• The SQL console</li> </ul>
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"> <li>• <a href="#">Systems</a>, which is the central access point for performing system administration and monitoring activities</li> <li>• <a href="#">Error Log</a>, which contains error and information messages</li> <li>• <a href="#">Properties</a>, which shows the detailed properties of the active resource (for example, the SAP HANA system selected in the <a href="#">Systems</a> view)</li> </ul> <p>To open a view, from the main menu, choose <b>▶ Window &gt; Show View ▾</b>.</p>

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

### 3.2.1.2 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
<i>Systems</i>	The <i>Systems</i> 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 <i>Systems</i> 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 <b>Window &gt; Show View &gt; Systems</b> .
<i>System Monitor</i>	The <i>System Monitor</i> is an editor that provides you with an overview of all your SAP HANA systems at a glance. From the <i>System Monitor</i> , you can drill down into the details of an individual system in the <i>Administration</i> editor.	In the toolbar of the <i>Systems</i> view, choose the  button.
<i>Administration</i>	The Administration editor the main tool for performing administration and monitoring activities.	You can access the <i>Administration</i> editor in several ways: <ul style="list-style-type: none"> <li>From the <i>Systems</i> view toolbar, choose the  <i>Open Default Administration</i> button.</li> <li>In the <i>Systems</i> view, double-click the system.</li> <li>In the context menu of the <i>Systems</i> view, choose <b>Configuration and Monitoring &gt; Open Administration</b>.</li> </ul>
<i>Administration Diagnosis Mode</i>	The <i>Administration</i> 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 <i>Administration</i> editor opens automatically in diagnosis mode in the following situations: <ul style="list-style-type: none"> <li>When you open the <i>Administration</i> editor for a system that cannot be reached by SQL</li> <li>When you initiate the start, stop, or restart of a system</li> </ul> You can also open the <i>Administration</i> editor in diagnosis mode from the <i>Systems</i> view toolbar by choosing the  <i>Open Diagnosis Mode</i> button.
<i>Backup</i>	The <i>Backup</i> editor is the main tool for performing administration and monitoring activities related to backup.	You can access the <i>Backup</i> editor in several ways: <ul style="list-style-type: none"> <li>Expand the system in the <i>Systems</i> view and choose the  <i>Backup</i> entry</li> <li>In the context menu of the <i>Systems</i> view, choose <b>Backup and Recovery &gt; Open Backup Console</b>.</li> </ul>

View/Editor	Description	How to Open
<a href="#">Security</a>	<p>The <a href="#">Security</a> editor is the main tool for managing the following aspects of security administration:</p> <ul style="list-style-type: none"> <li>• Password policy</li> <li>• Auditing</li> <li>• Data volume encryption</li> </ul>	<p>You can access the <a href="#">Security</a> editor in several ways:</p> <ul style="list-style-type: none"> <li>• Expand the system in the <a href="#">Systems</a> view and choose the  <a href="#">Security</a> entry</li> <li>• In the context menu of the <a href="#">Systems</a> view, choose  <a href="#">Security</a>  <a href="#">Open Security Console</a> .</li> </ul>
<a href="#">SQL Console</a>	<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"> <li>• From the <a href="#">Systems</a> view toolbar, choose the  button.</li> <li>• In the context menu of the <a href="#">Systems</a> view, choose <a href="#">Open SQL Console</a>.</li> </ul>

## Related Information

[Systems View \[page 67\]](#)

[System Monitor \[page 228\]](#)

[Administration Editor \[page 229\]](#)

[Backup Console \[page 917\]](#)

[Execute SQL Statements in SAP HANA Studio \[page 65\]](#)

## 3.2.2 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.

### **i** 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 \[page 70\]](#)

[Systems View \[page 67\]](#)

[SAP HANA Studio Administration Preferences \[page 86\]](#)

[User Authorization \[page 669\]](#)

## 3.2.3 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.

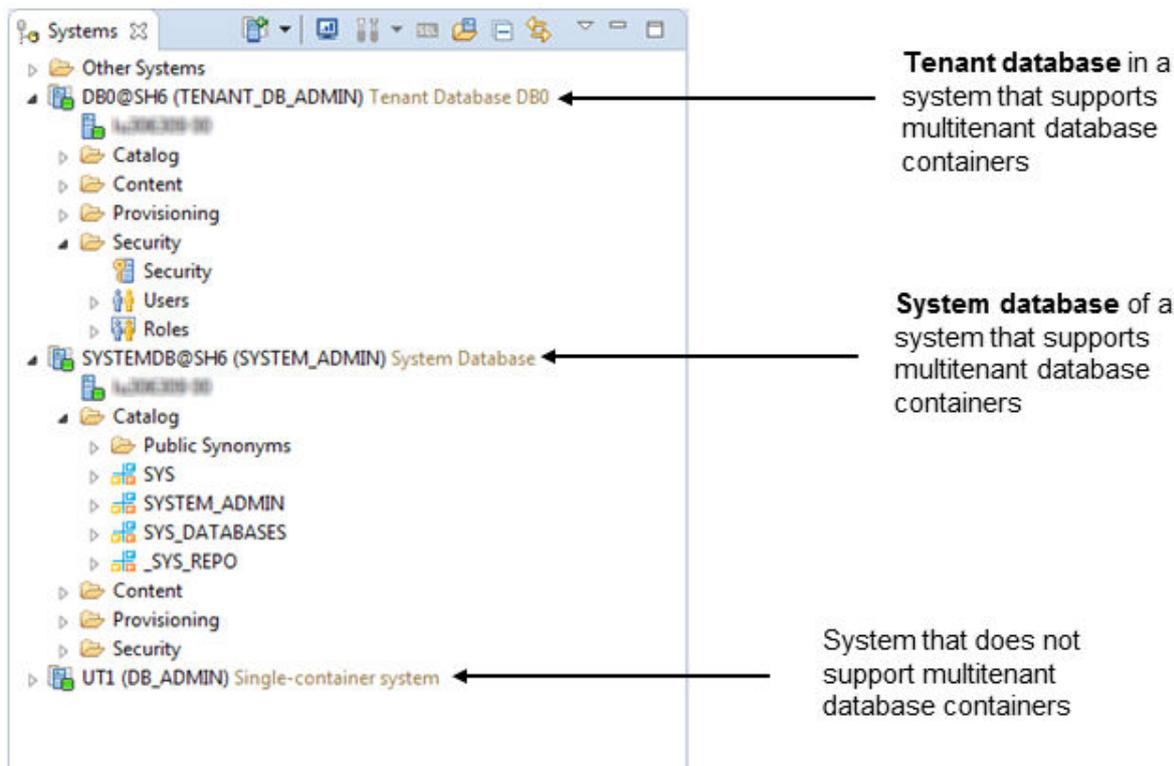
### 3.2.3.1 Systems View

The *Systems* view provides you with a hierarchical view of all the SAP HANA systems and SAP HANA multitenant database containers 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 and multitenant database containers. In the *Systems* view, systems that do not support multitenant database containers are identified solely by their SID. Multitenant database containers, or tenant databases, are identified by their database name and SID (<database\_name>@<SID>). The system database of a multiple-container system 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



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 \[page 69\]](#)

[Add an SAP HANA System \[page 70\]](#)

### 3.2.3.1.1 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
	<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 \[page 70\]](#)

[Link a Centrally-Stored Archive of SAP HANA Systems \[page 76\]](#)

[System Monitor \[page 228\]](#)

[Administration Editor \[page 229\]](#)

[Troubleshooting an Inaccessible or Unresponsive SAP HANA System \[page 481\]](#)

[Execute SQL Statements in SAP HANA Studio \[page 65\]](#)

## 3.2.3.2 Add an SAP HANA System

To work with and manage an SAP HANA system or multitenant database container 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
<b>Option</b>	JDBC connection parameter(s)
<b>Auto-Reconnect</b>	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 \[page 648\]](#)

[Provisioning Users \[page 702\]](#)

[System Connection Properties \[page 73\]](#)

[Configure System Usage Type \[page 221\]](#)

[Configure SSL for SAP HANA Studio Connections \[page 84\]](#)

[Creating and Configuring Tenant Databases \[page 104\]](#)

### 3.2.3.2.1 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"> <li> <b>Multi-host system</b>            If you are adding a multi-host system, specify the <b>master host</b>. 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.         </li> </ul> <div data-bbox="647 786 1394 994" style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;"> <p><b>→ Tip</b></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="647 1010 1394 1182" style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;"> <p><b>i Note</b></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"> <li> <b>Multitenant database containers</b>            If you are adding a tenant database in a multitenant system, specify the FQDN of the system host.         </li> </ul> <div data-bbox="647 1301 1394 1771" style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;"> <p><b>i Note</b></p> <p>To be able to access the SAP HANA cockpit from the SAP HANA studio, you can also enter the alias DNS name of the tenant database. Every tenant database requires an alias name so that HTTP requests can be dispatched to the XS server of the correct database. This is handled by the system's internal Web Dispatcher. Alias names are configured in the Web Dispatcher configuration file.</p> <p>If you don't enter the alias name of the tenant database here, you need to specify it explicitly as the XS server host in the system properties. You can do this after you have finished adding the system (▶ <i>Properties</i> ▶ <i>XS Properties</i> ▶).</p> </div>
Instance	Instance number of the system

Property	Description
Mode	<p>Whether the system you are connecting to is a <b>single-container system</b> or <b>multiple-container system</b></p> <p>If you are connecting to a system with multitenant database containers, then you must further specify the specific database that you want to connect to:</p> <ul style="list-style-type: none"> <li>• The system database</li> <li>• A tenant database</li> </ul>

### 3.2.3.3 Add an SAP HANA System with a Different User

If you want to work with an SAP HANA system or a multitenant database container 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
<b>Option</b>	JDBC connection parameter(s)
<b>Auto-Reconnect</b>	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 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 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 \[page 648\]](#)

[Provisioning Users \[page 702\]](#)

[Configure System Usage Type \[page 221\]](#)

[Configure SSL for SAP HANA Studio Connections \[page 84\]](#)

### 3.2.3.4 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 \[page 80\]](#)

### 3.2.3.5 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 \[page 78\]](#)

### 3.2.3.5.1 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.

### 3.2.3.6 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.

### 3.2.3.7 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.

---

## 3.2.3.8 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.

## 3.2.3.9 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 \[page 76\]](#)

### 3.2.3.10 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.

## 3.2.3.11 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.

#### 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.

## 3.2.3.12 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
<b>-h</b>	Host name
<b>-n</b>	Instance number
<b>-u</b>	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.

### 3.2.3.13 Configure 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`.

### **i** 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 \[page 70\]](#)

[Managing Client Certificates in the SAP HANA Database \[page 758\]](#)

## 3.2.4 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 \[page 86\]](#)
- [Global Settings \[page 87\]](#)
- [Runtime \[page 87\]](#)
- [Table Viewer \[page 90\]](#)

### 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 <a href="#">Systems</a> 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 <a href="#">Console</a> is available in the Administration editor. You can execute <code>hdbcons</code> commands directly in this console.</p> <div style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;"> <p> <b>Caution</b></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 <a href="#">Systems</a> view such as <a href="#">Tables</a> and <a href="#">Views</a> are opened.</p> <p>If you select this option, <b>all</b> 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 <a href="#">Fetch all database catalog objects</a> 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 <a href="#">Object limit reached</a> 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 <a href="#">Systems</a> 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 <ul style="list-style-type: none"> <li>Save content when SAP HANA Studio is closed</li> <li>Content save interval ... minutes</li> </ul>	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.  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: <ul style="list-style-type: none"> <li>Data separator</li> <li>Tab separated</li> <li>Align copied values with space</li> <li>Copy cell in editor by using <code>[CTRL] C</code></li> <li>Copy editor content with column header</li> </ul>	These are formatting options for copying content from the table editor.
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: <ul style="list-style-type: none"> <li>Show content</li> <li>Show definition</li> </ul>	This setting specifies which view of a table is opened when it is double-clicked in the <a href="#">Systems</a> view: its definition or its content.
Table Distribution Editor Maximum Number of Tables Displayed	This setting specifies the maximum number of tables that are displayed when you show table distribution.

#### 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 <a href="#">Result</a> 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 <a href="#">Result</a> 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 in 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"> <li>• Auto-commit mode</li> <li>• Isolation level</li> <li>• Confirm change of connection</li> </ul>	<ul style="list-style-type: none"> <li>• Auto-commit mode:               <ul style="list-style-type: none"> <li>◦ If on, the system performs all COMMIT actions automatically</li> <li>◦ If off, you have to enter COMMIT statements explicitly.</li> </ul> </li> <li>• The isolation level determines how the system implicitly controls locking and versioning of database objects.</li> <li>• 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.</li> </ul>
<b>SQL console settings</b>	
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	

---

## 4 System Administration

As a database administrator you are responsible for operational tasks related to the administration, monitoring, and maintenance of your SAP HANA systems.

### 4.1 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 studio and the SAP HANA cockpit. The SAP start service (`sapstartsrv`) is the standard SAP mechanism for starting and stopping systems.

#### Related Information

[Starting and Stopping Systems in SAP HANA Studio \[page 97\]](#)

[Starting and Stopping Systems in SAP HANA Cockpit \[page 91\]](#)

#### 4.1.1 Starting and Stopping Systems in SAP HANA Cockpit

Use the SAP HANA cockpit for offline administration to stop, start, or restart an SAP HANA system.

#### Related Information

[Open SAP HANA Cockpit for Offline Administration \[page 58\]](#)

## 4.1.1.1 Start a System

Use the SAP HANA cockpit for offline administration to start an SAP HANA system. In a system with multitenant database containers, all tenant databases will be started.

### 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 SAP HANA cockpit for offline administration.

You can do this the following ways:

Option	Description
<b>Directly</b>	<p>Enter the URL in your browser:</p> <pre>https://&lt;host&gt;:1129/lmsl/hdbcockpit/&lt;sid&gt;/index.html</pre> <p><b>i Note</b></p> <p>It's also possible to access the SAP HANA cockpit for offline administration via the URL <code>http://&lt;host&gt;:1128/lmsl/hdbcockpit/&lt;sid&gt;/index.html</code>). However, this is not recommended because passwords are transferred in plain text via HTTP.</p>
<b>From the SAP HANA cockpit</b>	<ol style="list-style-type: none"><li>1. Open the SAP HANA cockpit by entering the URL in your browser: <code>https://&lt;host&gt;:43&lt;instance&gt;/sap/hana/admin/cockpit</code> (recommended) or <code>http://&lt;host&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code></li><li>2. In the <i>SAP HANA Database Administration</i> group, click the tile <i>SAP HANA Cockpit for Offline Administration</i>.</li></ol> <p><b>i Note</b></p> <p>If you access the SAP HANA cockpit via HTTP, then the SAP HANA cockpit for offline administration is also accessed via HTTP. Therefore, we recommend configuring the SAP HANA cockpit for HTTPS access.</p>

2. Open the *System Operations* app by clicking the *Start, Stop, Restart* tile on the homepage of the SAP HANA cockpit for offline administration.
3. In the footer bar, choose *Start System*.

## Results

The database services start one by one. If the system contains multitenant database containers, the services of all tenant databases are started. For more information about how to stop an individual tenant database, see *Stop and Start a Tenant Database* in the *SAP HANA Administration Guide*.

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.

## Related Information

[View Diagnosis Files in SAP HANA Cockpit \[page 463\]](#)

### 4.1.1.2 Stop a System

Use the SAP HANA cockpit for offline administration to stop an SAP HANA system. In a system with multitenant database containers, all tenant databases will be stopped.

## 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 SAP HANA cockpit for offline administration.

You can do this the following ways:

Option	Description
Directly	Enter the URL in your browser: <code>https://&lt;host&gt;:1129/lmsl/hdbcockpit/&lt;sid&gt;/index.html</code>

Option	Description
	<p><b>i Note</b></p> <p>It's also possible to access the SAP HANA cockpit for offline administration via the URL <code>http://&lt;host&gt;:1128/lmsl/hdbcockpit/&lt;sid&gt;/index.html</code>). However, this is not recommended because passwords are transferred in plain text via HTTP.</p>
<b>From the SAP HANA cockpit</b>	<ol style="list-style-type: none"> <li>1. Open the SAP HANA cockpit by entering the URL in your browser: <code>https://&lt;host&gt;:43&lt;instance&gt;/sap/hana/admin/cockpit</code> (recommended) or <code>http://&lt;host&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code></li> <li>2. In the <i>SAP HANA Database Administration</i> group, click the tile <i>SAP HANA Cockpit for Offline Administration</i>.</li> </ol> <p><b>i Note</b></p> <p>If you access the SAP HANA cockpit via HTTP, then the SAP HANA cockpit for offline administration is also accessed via HTTP. Therefore, we recommend configuring the SAP HANA cockpit for HTTPS access.</p>

2. Open the *System Operations* app by clicking the *Start, Stop, Restart* tile on the homepage of the SAP HANA cockpit for offline administration.
3. In the footer bar, choose *Stop System*.
4. Specify how you want to stop the system:

Option	Description
<b>Softly</b>	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.
<b>Immediately</b>	The system is stopped immediately. Open transactions are aborted and rolled back.

## Results

The database services stop one by one. If the system contains multitenant database containers, the services of tenant databases are stopped. For more information about how to stop an individual tenant database, see *Stop and Start 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

[View Diagnosis Files in SAP HANA Cockpit \[page 463\]](#)

### 4.1.1.3 Restart a System

Use the SAP HANA cockpit for offline administration to restart an SAP HANA system. In a system with multitenant database containers, all tenant databases will be restarted.

#### 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 SAP HANA cockpit for offline administration.

You can do this the following ways:

Option	Description
<b>Directly</b>	<p>Enter the URL in your browser:</p> <pre>https://&lt;host&gt;:1129/lmsl/hdbcockpit/&lt;sid&gt;/index.html</pre> <p><b>i Note</b></p> <p>It's also possible to access the SAP HANA cockpit for offline administration via the URL <code>http://&lt;host&gt;:1128/lmsl/hdbcockpit/&lt;sid&gt;/index.html</code>). However, this is not recommended because passwords are transferred in plain text via HTTP.</p>
<b>From the SAP HANA cockpit</b>	<ol style="list-style-type: none"><li>1. Open the SAP HANA cockpit by entering the URL in your browser: <code>https://&lt;host&gt;:43&lt;instance&gt;/sap/hana/admin/cockpit</code> (recommended) or <code>http://&lt;host&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code></li><li>2. In the <i>SAP HANA Database Administration</i> group, click the tile <i>SAP HANA Cockpit for Offline Administration</i>.</li></ol>

Option	Description
	<p><b>i Note</b></p> <p>If you access the SAP HANA cockpit via HTTP, then the SAP HANA cockpit for offline administration is also accessed via HTTP. Therefore, we recommend configuring the SAP HANA cockpit for HTTPS access.</p>

2. Open the *System Operations* app by clicking the *Start, Stop, Restart* tile on the homepage of the SAP HANA cockpit for offline administration.
3. In the footer bar, choose *Restart System*.
4. Specify how you want to stop the system:

Option	Description
<b>Softly</b>	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.
<b>Immediately</b>	The system is stopped immediately. Open transactions are aborted and rolled back.

## Results

The database services first stop one by one and then restart one by one. If the system contains multitenant database containers, the services of all tenant databases are stopped and started. For more information about how to stop an individual tenant database, see *Stop and Start a Tenant Database* in the *SAP HANA Administration Guide*.

When all services have restarted, the system has the status *Running*.

### → Tip

To analyze any problems that may occur during start-up or while the system is stopped, you can access the system's diagnosis files from the homepage of the SAP HANA cockpit.

## Related Information

[View Diagnosis Files in SAP HANA Cockpit \[page 463\]](#)

## 4.1.2 Starting and Stopping Systems in SAP HANA Studio

Use the SAP HANA studio to stop, start, or restart an SAP HANA system.

### 4.1.2.1 Start a System

Use the SAP HANA studio to start an SAP HANA system. In a system with multitenant database containers, 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...* ►

#### Note

If the system supports multitenant database containers, 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.
- If the system contains multitenant database containers, 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 <sid>adm \[page 647\]](#)

[Stop and Start a Tenant Database \[page 134\]](#)

[Monitoring System Availability \[page 230\]](#)

[Starting and Stopping Distributed SAP HANA Systems Using SAPControl \[page 1004\]](#)

### 4.1.2.2 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

If the system supports multitenant database containers, 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
<b>Soft</b>	The system is stopped after all running statements have finished or the specified timeout is reached.
<b>Hard</b>	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.
- If the system contains multitenant database containers, 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 `<sid>adm` \[page 647\]](#)

[Monitoring System Availability \[page 230\]](#)

[Stop and Start a Tenant Database \[page 134\]](#)

[Starting and Stopping Distributed SAP HANA Systems Using SAPControl \[page 1004\]](#)

### 4.1.2.3 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

If the system supports multitenant database containers, 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
<b>Soft</b>	The system is stopped after all running statements have finished or the specified timeout is reached.
<b>Hard</b>	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.
- If the system contains multitenant database containers, 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 \[page 101\]](#)

[Operating System User <sid>adm \[page 647\]](#)

[Stop and Start a Tenant Database \[page 134\]](#)

## 4.1.2.3.1 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.

## 4.1.2.4 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
<b>Stop...</b>	The service is stopped normally and then typically restarted.
<b>Kill...</b>	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.
<b>Reconfigure Service...</b>	The service is reconfigured. This means that any changes made to parameters in the system's configuration files are applied.
<b>Start Missing Services...</b>	Any inactive services are started.

---

## Related Information

[Monitoring Status and Resource Usage of System Components \[page 234\]](#)

[High Availability for SAP HANA \[page 774\]](#)

[View Diagnosis Files in SAP HANA Studio \[page 461\]](#)

### 4.1.2.5 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 `<sid>adm` \[page 647\]](#)

[Monitoring Overall System Status and Resource Usage \[page 232\]](#)

[Troubleshooting an Inaccessible or Unresponsive SAP HANA System \[page 481\]](#)

## 4.2 Managing Multitenant Database Containers

As the administrator of a multiple-container 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

Basic administration of tenant databases is possible using the SAP HANA cockpit. However, the SAP HANA studio and command-line tools are required for some tasks.

### **i** Note

If you have SAP HANA options installed, review the section about multitenant database containers 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 \[page 104\]](#)

[Monitoring and Managing Tenant Databases \[page 129\]](#)

[Managing Resources in Multiple-Container Systems \[page 143\]](#)

[Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities \[page 1360\]](#)

### 4.2.1 Creating and Configuring Tenant Databases

During the installation of a multiple-container system, only the system database is initially created. You create and configure tenant databases afterward. If your system was updated or installed in single-container mode, you must convert it to multiple-container mode before setting up tenant databases.

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)
- 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

## Note

Basic administration of tenant databases is possible using the SAP HANA cockpit. However, the SAP HANA studio and command-line tools are required for some tasks.

## Related Information

[Converting an SAP HANA System to Support Multitenant Database Containers \[page 572\]](#)

[Increase the System Isolation Level \[page 105\]](#)

[Create a Tenant Database \[page 111\]](#)

[Disable Features on a Tenant Database \[page 114\]](#)

[Enable and Configure Cross-Database Access \[page 115\]](#)

[Prevent Changes to System Properties in Tenant Databases \[page 121\]](#)

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

[Administration of Multitenant Database Containers \[page 19\]](#)

### 4.2.1.1 Increase the System Isolation Level

You can increase the isolation level of an existing multiple-container 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 can log on to the system database and you have the system privilege `DATABASE ADMIN`.
- Internal SAP HANA communication has been appropriately configured for SSL (secure sockets layer). The property `[communication] ssl` in the `global.ini` file must have the value `false` (default) or `systempki`.

### Caution

The value of the `[communication] ssl` property must **not** be `true`. This will be the case if you are using a manually configured public key infrastructure (PKI) to secure internal communication between hosts. You can switch to system PKI by changing setting the parameter to `systempki`

For more information, see *Secure 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.

You do this in the system database in one of the following ways:

- Executing the SQL statement `ALTER SYSTEM STOP DATABASE <databasename>`
- Using 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 <databasename> OS USER '<username>' OS GROUP '<groupname>'`

5. Start all tenant databases.

You do this in the system database in one of the following ways:

- Executing the SQL statement `ALTER SYSTEM START DATABASE <database_name>`
- Using 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, see *Securing Internal Communication* in the *SAP HANA Security Guide*.
- Operations that require operating system access are restricted to users with the correct permissions. For more information, see *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 \[page 107\]](#)

[Stop and Start a Tenant Database \[page 134\]](#)

[Create a Tenant Database \[page 111\]](#)

[Isolation Level High for Backups With SAP HANA Multitenant Database Containers and Third-Party Tools \[page 913\]](#)

### 4.2.1.1.1 Database Isolation

Every tenant database in a multiple-container system 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 in a multiple-container system 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

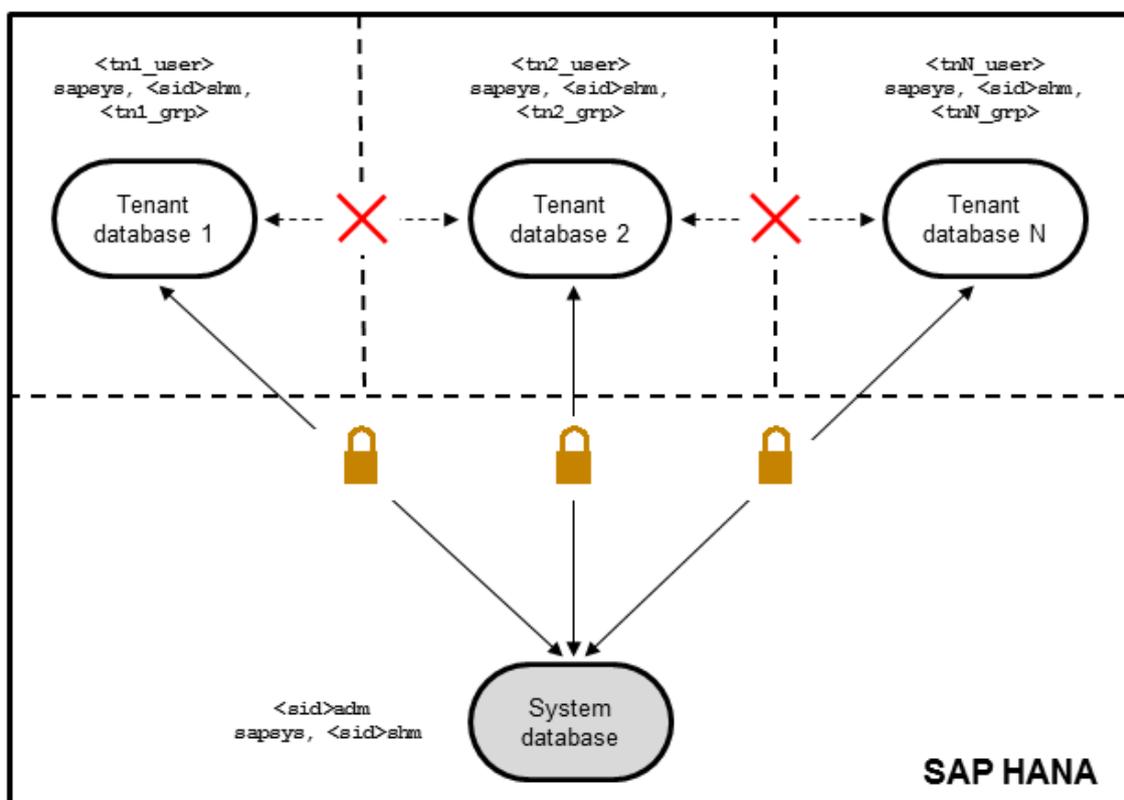
`<sid>adm` is the OS user for the system database.

## 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. 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.

For more information, see:

- *Installing a Multitenant Database Container SAP HANA System* in the *SAP HANA Server Installation and Update Guide*
- *Increase the System Isolation Level* in the *SAP HANA Administration Guide*
- *Secure Internal Communication* in the *SAP HANA Security Guide*

## Related Information

[Increase the System Isolation Level \[page 105\]](#)

[File and Directory Permissions with High Isolation \[page 109\]](#)

### 4.2.1.1.2 File and Directory Permissions with High Isolation

In a multiple-container system configured for high isolation, database-specific data on the file system is protected using standard file and directory permissions.

The following table shows who has access to which data on the file system:

File and Directory Permissions	Tenant OS User in Tenant OS Group	<sid>adm User
Read, write, and execute permissions (0770) on database-specific directories containing: <ul style="list-style-type: none"> <li>• Data volumes</li> <li>• Log volumes</li> <li>• Log mirror volumes</li> <li>• Backups</li> </ul>	Yes	No
Read, write, and execute permissions (0770) on database-specific directories containing: <ul style="list-style-type: none"> <li>• Configuration (*.ini) files</li> <li>• Trace files</li> </ul>	Yes	Yes
Read and write permissions (0666) on files in database-specific directory containing: <ul style="list-style-type: none"> <li>• Configuration (*.ini) files</li> <li>• Trace files</li> </ul>	Yes	Yes
Read permission (644) on files in directory containing system configuration files	Yes	Yes
Read and write permissions (600) on files in trace directory of the system database	No	Yes

## 4.2.1.2 Decrease the System Isolation Level

If you configured a multiple-container 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 can log on to the system database and you have the system privilege DATABASE ADMIN.

### Procedure

1. Stop all tenant databases in the system.

You do this in the system database in one of the following ways:

- Executing the SQL statement `ALTER SYSTEM STOP DATABASE <database_name>`
- Using 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.

4. Start all tenant databases.

You do this in the system database in one of the following ways:

- Executing the SQL statement `ALTER SYSTEM START DATABASE <databasename>`
- Using 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

[Stop and Start a Tenant Database \[page 134\]](#)

### 4.2.1.3 Create a Tenant Database

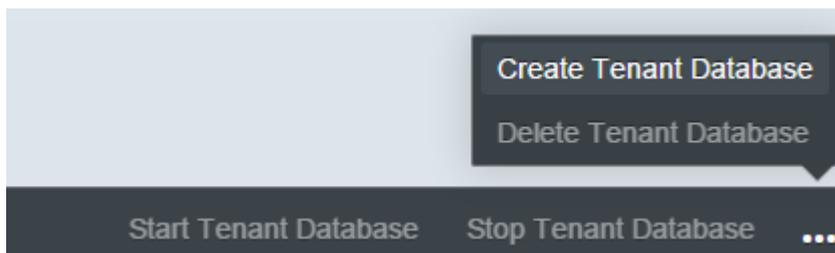
You create tenant databases after installation of a multiple-container system, after conversion from a single-container system to a multiple-container system, or anytime a new database is needed. You create tenant databases from the system database using the *Manage Databases* app of the SAP HANA cockpit.

#### Prerequisites

- You are connected to the system database.
- You have the privileges granted by role `sap.hana.admin.cockpit.sysdb.roles::SysDBAdmin`. You can grant roles using the *Assign Roles* app of the SAP HANA cockpit. For more information, see *Assign Roles to a User* in the *SAP HANA Administration Guide*.
- The *Manage Databases* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA System Administration* tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.
- 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 *Database Isolation*.

#### Procedure

1. Open the *Manage Databases* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.
2. In the footer toolbar, open the overflow menu and choose *Create Tenant Database*.



The *Create Tenant Database* app 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 Multitenant Database Containers* 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](#).

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 SQL port number and HTTP port number of the embedded XS server are assigned. These are the internal communication port number plus 1 and 2 respectively.

- 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 the [Manage Databases](#) app. 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 multitenant database containers in the *SAP HANA Administration Guide*.

## Related Information

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Database Isolation \[page 107\]](#)

[Linux Kernel Parameters \[page 528\]](#)

[Creating Backups \[page 920\]](#)

[Disable Features on a Tenant Database \[page 114\]](#)

[Enable and Configure Cross-Database Access \[page 115\]](#)

[Prevent Changes to System Properties in Tenant Databases \[page 121\]](#)

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

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

## 4.2.1.4 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 studio.

### Prerequisites

- You are logged on to the system database.
- The tenant database exists.
- 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 Multitenant Database Containers* 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. In the Administration editor, choose the *Configuration* tab and navigate to the `global.ini` file.
3. Disable a feature as follows:
  - a. Navigate to the `customizable_functionalities` section.
  - b. Right-click the property that corresponds to the feature you want to disable and choose *Change*.

- c. Select the relevant tenant database(s) and enter **false** as the new value.

#### **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 value.

- d. Choose *Save*.
4. Disable further features not required in the tenant database(s).
5. Restart the affected tenant database(s).

## Results

The feature is disabled. You can verify this in the view `M_CUSTOMIZABLE_FUNCTIONALITIES (SYS_DATABSES)`.

Tenant database administrators can see which features are enabled in their database using the view `M_CUSTOMIZABLE_FUNCTIONALITIES (SYS)`.

## Related Information

[Stop and Start a Tenant Database \[page 134\]](#)

[Monitoring Using System and Statistics Views \[page 283\]](#)

### 4.2.1.5 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 in a multiple-container system 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 `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**.

### 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 \[page 117\]](#)

[Troubleshooting Error Situations Related to Cross-Database Access \[page 118\]](#)

### 4.2.1.5.1 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 in a multiple-container system 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 *Multitenant Database Containers 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 \[page 115\]](#)

[Troubleshooting Error Situations Related to Cross-Database Access \[page 118\]](#)

[Workload Management and Cross-Database Queries \[page 118\]](#)

### 4.2.1.5.2 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.

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 \[page 418\]](#)

### 4.2.1.5.3 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?

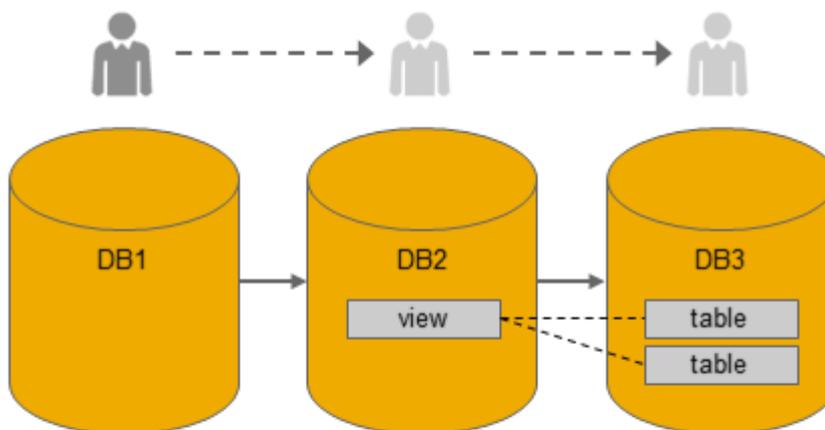
Cross-database access supports only read-only operations. Changes to an object on one tenant database cannot therefore be reflected accurately on other tenant databases that contain objects dependent on the changed object. This affects the validity flag in the relevant system views, as well as the object dependencies. Remote objects may stay valid if they retain their internal object identifier during re-creation and are re-created in a 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?

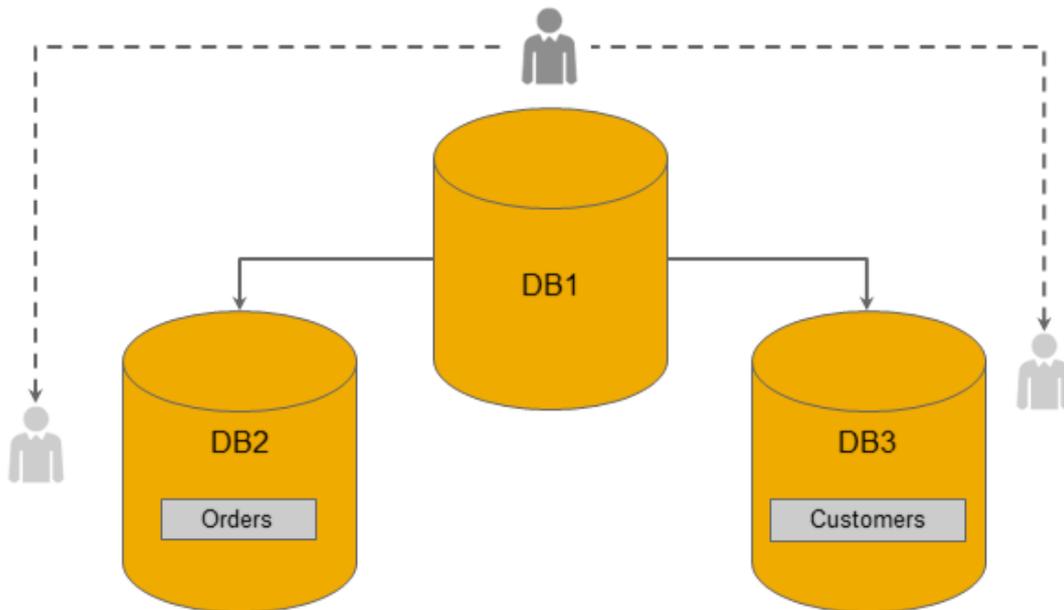
Cross-database queries do not support multiple tenant database levels as part of a view hierarchy, even if communication between databases is enabled (including the required authorized remote users).

### 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;
```

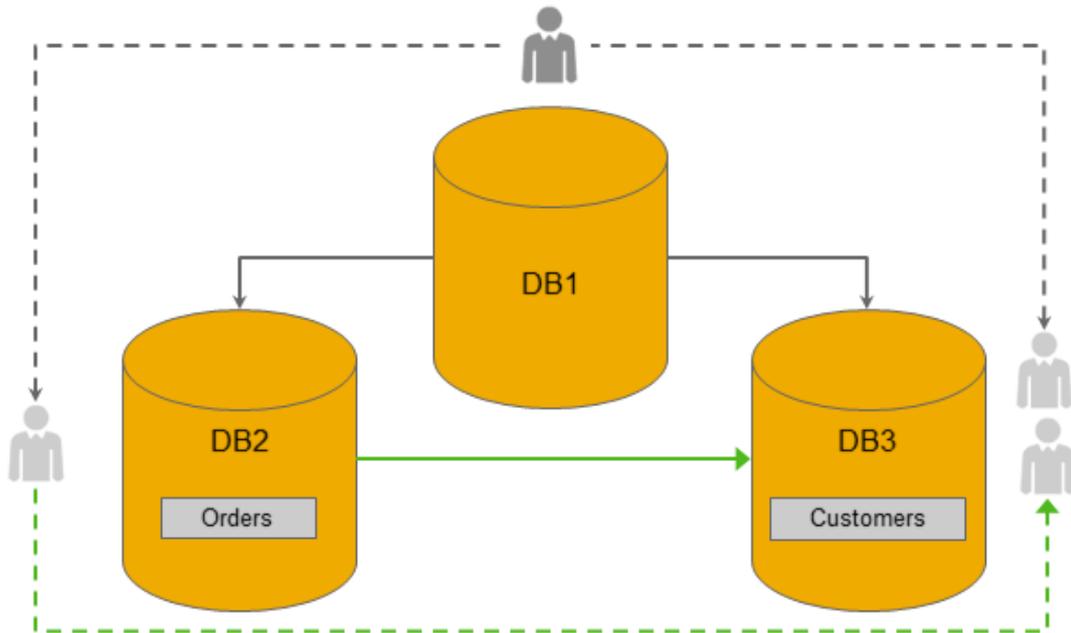


### What's the problem?

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 \[page 115\]](#)

### 4.2.1.6 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 studio.

## Prerequisites

- You are logged on to the system database.
- 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 studio.

### **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. In the Administration editor, choose the *Configuration* tab and navigate to the `multidb.ini` file.
2. Add new properties to the blacklist as follows:
  - a. Right-click the `readonly_parameters` section and choose *Add Parameter*.
  - b. 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.
Host	Configuration not possible in any tenant database on the specified host(s).
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.

- c. Choose *Next*.

- d. In the *Key* field, enter the 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`.
- e. 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/*`).
- f. Choose *Finish*.
- g. Add further properties to the blacklist 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 \(INI Files\) \[page 212\]](#)

## 4.2.1.6.1 Default Blacklisted System Properties in Multitenant Database Containers

In systems that support multitenant database containers, 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, see *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"> <li>default_audit_trail_type</li> <li>emergency_audit_trail_type</li> <li>alert_audit_trail_type</li> <li>critical_audit_trail_type</li> </ul>	Prevents configuration of audit trail targets
indexserver.ini/ authentication	SapLogonTicketTrustStore	Prevents configuration of the trust store for user authentication with logon/assertion tickets
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 option)
global.ini/persistence	<ul style="list-style-type: none"> <li>basepath_datavolumes_es</li> <li>basepath_logvolumes_es</li> <li>basepath_databackup_es</li> <li>basepath_logbackup_es</li> </ul>	
multidb.ini/ readonly_parameters	*	Prevents configuration of the <code>multidb.ini</code> file itself
memorymanager	<ul style="list-style-type: none"> <li>allocationlimit</li> <li>minallocationlimit</li> <li>global_allocation_limit</li> <li>async_free_threshold</li> <li>async_free_target</li> </ul>	Prevents configuration of memory allocation parameters

File/Section	Properties	Description
execution	max_concurrency	Prevents configuration of threading and parallelization parameters
session	<ul style="list-style-type: none"> <li>maximum_connections</li> <li>maximum_external_connections</li> </ul>	
sql	sql_executors	

## Related Information

[Prevent Changes to System Properties in Tenant Databases \[page 121\]](#)

### 4.2.1.7 Configure HTTP(S) Access to Multitenant Database Containers

To enable Web-based applications to send HTTP(S) requests to multitenant database containers, 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 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 in a multiple-container system, 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** or is not available (revisions earlier than SPS 10), 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 (as of SPS 11)  
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"> <li>1. Open the Administration editor and choose the <i>Configuration</i> tab.</li> <li>2. Navigate to the <code>xsengine.ini</code> file and expand the <code>public_urls</code> section.</li> <li>3. For each tenant database in the system, add the new properties <code>http_url</code> and <code>https_url</code> at the <b>database layer</b> and enter its public URL as the value: <code>http://&lt;virtual_hostname&gt;:80&lt;instance&gt;</code></li> </ol>
SQL	For each tenant database, execute the statements: <ul style="list-style-type: none"> <li>◦ ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', '&lt;tenant_DB_name&gt;') SET ('public_urls', 'http_url') = 'http://&lt;virtual_hostname&gt;:80&lt;instance&gt;' WITH RECONFIGURE;</li> </ul>

Option	Description
	<ul style="list-style-type: none"> <li>ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', '&lt;tenant_DB_name&gt;') SET ('public_urls', 'https_url') = 'https://&lt;virtual_hostname&gt;:43&lt;instance&gt;' WITH RECONFIGURE;</li> </ul>

### **i** 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
<b>Single certificate for all databases</b>	<ol style="list-style-type: none"> <li>Start the SAP HANA Web Dispatcher Administration tool at <code>http://&lt;localhost&gt;:80&lt;instance&gt;/sap/hana/xs/wdisp/admin/</code>.</li> <li>For the default <code>SAPSSLS.pse</code> trust store, create a new SSL key pair and certificate request: <ol style="list-style-type: none"> <li>From the main menu, choose <b>SSL and Trust Configuration</b> &gt; <b>PSE Management</b>.</li> <li>From the <b>Manage PSE</b> menu, choose <code>SAPSSLS.pse</code>.</li> <li>Choose <b>Recreate PSE</b>.</li> <li>Enter a distinguished name that matches the host name of all tenant databases.</li> </ol> </li> </ol> <div style="background-color: #fff9c4; padding: 10px; margin-top: 10px;"> <p><b>Example</b></p> <ul style="list-style-type: none"> <li>Physical host name: <code>myhost.example.com</code></li> <li>Tenant host name 1: <code>mydatabase1.example.com</code></li> <li>Tenant host name 2: <code>mydatabase2.example.com</code></li> </ul> </div>

Option	Description
	<p>In this case, you specify <b>CN=*.example.com</b> as the DN, thus creating a server certificate that matches all tenant databases and the system database.</p> <ol style="list-style-type: none"> <li>Choose <a href="#">Create</a>.</li> <li>Create a certificate request and submit to your certificate authority (CA) for signing (<a href="#">Create CA Response</a>).</li> <li>Import the signed certificate</li> </ol> <p>For more information, see <i>Configure HTTPS (SSL) for Client Application Access</i>.</p>
<b>Individual certificates for each database</b>	<ol style="list-style-type: none"> <li>Start the SAP HANA Web Dispatcher Administration tool at <code>http://&lt;localhost&gt;:80&lt;instance&gt;/sap/hana/xs/wdisp/admin/</code>.</li> <li>For each tenant database and the system database, create a new trust store with a unique certificate: <ol style="list-style-type: none"> <li>From the main menu, choose <b>SSL and Trust Configuration</b> <b>PSE Management</b>.</li> <li>On the PSE management screen, choose <a href="#">Create New PSE</a>.</li> <li>Enter a file name for the new PSE.</li> </ol> <div data-bbox="480 853 1394 958" style="background-color: #fff9c4; padding: 5px;"> <p><b>Example</b> example.pse</p> </div> <ol style="list-style-type: none"> <li>Enter the distinguished name: <b>CN=&lt;host name used for the tenant database in the public_urls section of the xsengine.ini file&gt;</b></li> <li>Choose <a href="#">Create</a>.</li> <li>For the new PSE, create a certificate request and submit to your CA for signing (<a href="#">Create CA Response</a>).</li> <li>Import the signed certificate into the new PSE (<a href="#">Import CA Response</a>).</li> </ol> </li> <li>Configure the Web Dispatcher to use multiple certificates: <ol style="list-style-type: none"> <li>In the <code>webdispatcher.ini</code> file, create or change the parameter <code>[profile] icm/ssl_config_0</code>, specifying as the value: <b>ID=ssl_config_main, CRED=SAPSSLS.pse, SNI_CREDS=&lt;semicolon (';') separated list of database PSE files&gt;</b></li> <li>Add <b>,SSLCONFIG=ssl_config_main</b> to the value of the <code>icm/server_port</code> parameter for the HTTPS port (by default <code>icm/server_port_1</code>).</li> </ol> <div data-bbox="480 1435 1394 1579" style="background-color: #fff9c4; padding: 5px;"> <p><b>Example</b></p> <pre>icm/server_port_1 = PROT=HTTPS,PORT=4443\$ (SAPSYSTEM),PROCTIMEOUT=600, SSLCONFIG=ssl_config_main</pre> </div> </li> </ol>

## 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 \[page 1067\]](#)

[Using SAP Web Dispatcher for Load Balancing with Tenant Databases \[page 182\]](#)

## 4.2.2 Monitoring and Managing Tenant Databases

To ensure the overall performance and stability of a multiple-container 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.

### Note

Basic administration of tenant databases is possible using the SAP HANA cockpit. However, the SAP HANA studio and 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 Multitenant Database Containers \[page 19\]](#)

[Monitor Status and Resource Usage of Tenant Databases \[page 130\]](#)

[Monitor Alerts in Tenant Databases \[page 133\]](#)

[Stop and Start a Tenant Database \[page 134\]](#)

[Add a Service to a Tenant Database \[page 136\]](#)

[Remove a Service from a Tenant Database \[page 138\]](#)

[Delete a Tenant Database \[page 140\]](#)

[View Diagnosis Files of an Unavailable Tenant Database \[page 141\]](#)

### 4.2.2.1 Monitor Status and Resource Usage of Tenant Databases

You monitor the overall availability, resource usage, and performance of tenant databases from the system database using the *Manage Databases* app of the SAP HANA cockpit. From the *Manage Databases* app, you can then drill down to dedicated apps for more detailed information about individual tenant databases.

## Prerequisites

- You are connected to the system database.
- You have the privileges granted by role `sap.hana.admin.cockpit.sysdb.roles::SysDBAdmin`. You can grant roles using the *Assign Roles* app of the SAP HANA cockpit. For more information, see *Assign Roles to a User* in the *SAP HANA Administration Guide*.
- The *Manage Databases* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA System Administration* tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

## Procedure

Open the *Manage Databases* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.

You see the status of all the databases in the system, the number of medium and high priority alerts in each database, as well as information about the resource utilization of each database (memory, CPU, and disk).

## Next Steps

- If you want to examine a particular database in more detail, you can drill down further by clicking the aspect you're interested in (for example, database name, alerts or used memory). For more information about drill-down options, see *Database Details*.
- Depending on the situation, you may need to perform further operations on a tenant database (stop, start, and delete).

## Related Information

[Database Details \[page 131\]](#)

[Stop and Start a Tenant Database \[page 134\]](#)

[Delete a Tenant Database \[page 140\]](#)

### 4.2.2.1.1 Database Details

The *Manage Databases* app provides you with detailed information about all databases in a multiple-container system, 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
<i>Database Name</i>	Name of the tenant database	Click the database name to open the <i>Manage Services</i> app  The <i>Manage Services</i> app allows you to analyze the status and resource usage of the individual services of the database. For more information, see <i>Service Details</i> . From the <i>Manage Services</i> app, you can also stop and start services.

Column	Description	Drill-Down Option
<a href="#">Status</a>	Status of the tenant database: <ul style="list-style-type: none"> <li>• <a href="#">Running</a></li> <li>• <a href="#">Not running</a></li> <li>• <a href="#">Starting</a></li> <li>• <a href="#">Stopping</a></li> </ul>	No drill-down available
<a href="#">Alerts</a>	The number of high and medium priority alerts in the database	<p>Click the number of alerts to open them in the <a href="#">Alerts</a> app</p> <p>The <a href="#">Alerts</a> app allows you to view and analyze alerts occurring in the database. You can view past occurrences of alerts. For more information, see <a href="#">Alert Details</a> and <a href="#">Alert Priorities</a>. You can also see how alerts are configured. For more information, see <a href="#">Alert Checker Details</a> and <a href="#">Alert Checker Statuses</a>.</p> <div style="background-color: #fff9c4; padding: 10px;"> <p><b>i Note</b></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 <b>not</b> visible to the system administrator in the system database.</p> </div>
<a href="#">Start Time</a>	The time of the most recent start of the database	No drill-down available
<a href="#">Used Memory</a>	The used memory of the database in relation to the system	<p>Click the used memory bar to open the <a href="#">Performance Monitor</a> app</p> <p>The <a href="#">Performance Monitor</a> app allows you to visually analyze historical performance in the database across a range of related performance indicators. For more information, see <a href="#">Key Performance Indicators</a>.</p>
<a href="#">CPU Usage</a>	The CPU usage of the database in relation to the system	<p>Click the CPU usage bar to open the <a href="#">Performance Monitor</a> app</p> <p>The <a href="#">Performance Monitor</a> app allows you to visually analyze historical performance in the database across a range of related performance indicators. For more information, see <a href="#">Key Performance Indicators</a>.</p>

Column	Description	Drill-Down Option
<a href="#">Disk Usage</a>	The disk usage of the database in relation to the system	<p>Click the disk usage bar to open the <a href="#">Performance Monitor</a> app</p> <p>The <a href="#">Performance Monitor</a> app allows you to visually analyze historical performance in the database across a range of related performance indicators. For more information, see <a href="#">Key Performance Indicators</a>.</p>

## Related Information

[Service Details \[page 293\]](#)

[Alert Details \[page 302\]](#)

[Alert Priorities \[page 243\]](#)

[Alert Checker Details \[page 305\]](#)

[Alert Checker Statuses \[page 306\]](#)

[Key Performance Indicators \[page 314\]](#)

### 4.2.2.2 Monitor Alerts in Tenant Databases

Alert situations in tenant databases may potentially impact the health of the overall system. For this reason, you as system administrator can monitor certain alerts occurring in individual tenant databases. You can do this using the [Alerts](#) app of the SAP HANA cockpit.

## Prerequisites

- You are connected to the system database.
- You have the privileges granted by role `sap.hana.admin.cockpit.sysdb.roles::SysDBAdmin`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [System Alerts](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the [SAP HANA System Administration](#) tile catalog. For more information, see [Customizing the Homepage of SAP HANA Cockpit](#).

## Context

As the administrator of a multiple-container system, you can monitor certain alerts occurring in individual tenant databases using the [Alerts](#) app. You can see only those 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.

## Procedure

Open the *Alerts* app by clicking the *System Alerts* tile on the homepage of the SAP HANA cockpit.

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 in the *Alerts* app from 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, simply click *View Alert Configuration* in the footer toolbar. You cannot change the configuration.

## Related Information

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[System and Statistics Views in Multiple-Container Systems \[page 142\]](#)

[Alert Details \[page 302\]](#)

[Alert Priorities \[page 243\]](#)

[Alert Checker Details \[page 305\]](#)

[Alert Checker Statuses \[page 306\]](#)

### 4.2.2.3 Stop and Start a Tenant Database

You stop and start tenant databases from the system database using the *Manage Databases* app of the SAP HANA cockpit.

## Prerequisites

- You are connected to the system database.

- You have the privileges granted by role `sap.hana.admin.cockpit.sysdb.roles::SysDBAdmin`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- If you are stopping the database, consider backing it up first. For more information, see [Performing Backups](#).

## Context

As a system administrator, you can start tenant databases either individually, or all at once by starting the whole system. The same applies for stopping tenant databases.

### **i** Note

If you stop a tenant database individually, you can subsequently only start it again individually. It will not be started with a full system (re)start. For more information about full system stop, start, and restart see [Stop a System](#), [Start a System](#), and [Restart a System](#) in the *SAP HANA Administration Guide*.

## Procedure

1. Open the [Manage Databases](#) app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.
2. Select the tenant database that you want to stop or start.
3. Choose [Stop Tenant Database](#) or [Start Tenant Database](#) in the footer toolbar. The database starts stopping or starting. This may take a few moments.

## Results

- The database is started or stopped.

### **i** Note

If you stopped the database, it 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

[Start a System \[page 97\]](#)

---

[Stop a System \[page 98\]](#)

[Restart a System \[page 99\]](#)

[Creating Backups \[page 920\]](#)

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

## 4.2.2.4 Add a Service to a Tenant Database

To scale out a tenant database and/or distribute it across multiple hosts, you can add further server components, for example, an additional index server or a separate XS server. You add a service to a tenant database using the ALTER DATABASE SQL command.

### Prerequisites

- You are logged on to the system database.
- You have the system privilege DATABASE ADMIN.

### Context

You can add any service that persists data to an existing tenant database. You do this from the system database using the ALTER DATABASE ADD `<service>` statement. The main services that you may need to add are the following:

- indexserver
- xsengine

#### **i** 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. If you have already configured the SAP Web Dispatcher for HTTP(s) access based on the embedded xsengine service, you must update the configuration.

- scriptserver
- dpserver

#### **i** Note

You cannot add a statisticsserver service. This always runs embedded in the master index server of a tenant database.

In addition to the service you want to add, you may also specify the following information:

- The host on which the service is to be added

- The number of the internal communication port of the new service

**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. You do this by configuring 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.

**⚠ Caution**

Adding a service 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. Open the SQL console of the SAP HANA studio.
2. Execute the ALTER DATABASE statement.

For more information about the syntax of this statement, see *ALTER DATABASE* in the *SAP HANA SQL and System Views Reference*.

**🧩 Example**

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:30146'</code>
<p><b>i Note</b></p> <p>The SQL port number is the internal communication port number plus 1.</p>	
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>

---

## Results

- 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.

You can see the new service in the SAP HANA studio on the [Services](#) tab of the Administration editor and in the *Manage Services* app of the SAP HANA cockpit.

## Next Steps

- If you added an xsengine service, update the configuration of the SAP Web Dispatcher. For more information, see *Configure HTTP Access to Multitenant Database Containers*.
- Perform a full backup (data backup or storage snapshot). For more information, see *Perform a Complete Data Backup (SAP HANA Studio)*.

## Related Information

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 922\]](#)

[Monitoring Status and Resource Usage of System Components \[page 234\]](#)

[Execute SQL Statements in SAP HANA Studio \[page 65\]](#)

[The Statistics Service \[page 245\]](#)

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

## 4.2.2.5 Remove a Service from a Tenant Database

You can remove server components that are no longer needed from a tenant database. Removing services may be necessary, for example, if you are preparing to do a database copy using backup and recovery. You remove a service from a tenant database using the ALTER DATABASE SQL command.

## Prerequisites

- You are logged on to the system database.
- You have the system privilege DATABASE ADMIN.

## Context

You can remove any service from an existing tenant database, with the exception of the master indexserver service. You do this from the system database using the ALTER DATABASE REMOVE <service> statement.

You must specify the host and internal communication port number of the service to be removed.

### Caution

Removing a service 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 removing a service.

## Procedure

1. Open the SQL console of the SAP HANA studio.
2. Execute the ALTER DATABASE statement.

For more information about the syntax of this statement, see *ALTER DATABASE* in the *SAP HANA SQL and System Views Reference*.

### Example

```
ALTER DATABASE DB0 REMOVE 'indexserver' AT LOCATION 'HOST_A:30146'
```

### Note

The port number specified is that of the internal communication port of the service being removed.

## Results

- 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.

## Next Steps

- If you removed an xsengine service, update the configuration of the SAP Web Dispatcher. For more information, see *Configure HTTP Access to Multitenant Database Containers*.
- Perform a full backup (data backup or storage snapshot). For more information, see *Perform a Complete Data Backup (SAP HANA Studio)*.

## Related Information

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 922\]](#)

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

### 4.2.2.6 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

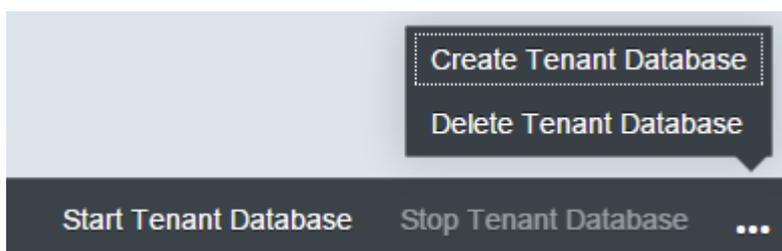
- You are connected to the system database.
- You have the privileges granted by role `sap.hana.admin.cockpit.sysdb.roles::SysDBAdmin`. You can grant roles using the *Assign Roles* app of the SAP HANA cockpit. For more information, see *Assign Roles to a User* in the *SAP HANA Administration Guide*.
- The *Manage Databases* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA System Administration* tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

#### Procedure

1. Open the *Manage Databases* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.
2. Stop the tenant database that you plan to delete by selecting it and then clicking *Stop Tenant Database* in the footer toolbar.

The system starts stopping the database. Once stopped, its status changes to *Not running*.

3. In the footer toolbar, open the overflow menu and choose *Delete Tenant Database*.



---

## Results

The system starts deleting the database. Once deleted, it disappears from the list of databases. Volumes, trace files, and file-based backups are removed. Backup directories that were previously in use, and backups that are written to by third-party backup tools, are not deleted.

## 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 in SAP HANA Studio \[page 65\]](#)

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

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

### 4.2.2.7 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 studio.

## Prerequisites

- You are logged on to the system database.
- You have the system privilege CATALOG READ.

## Procedure

1. In the SAP HANA studio, open the Administration editor and choose *Diagnosis Files*.  
The diagnosis files of the system database are displayed.
2. Using the *Database* filter, select the tenant database(s) whose diagnosis files you want to see.  
The diagnosis files of the selected tenant database(s) 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

[View Diagnosis Files in SAP HANA Studio \[page 461\]](#)

[Options for Diagnosis File Handling \(SAP HANA Studio\) \[page 461\]](#)

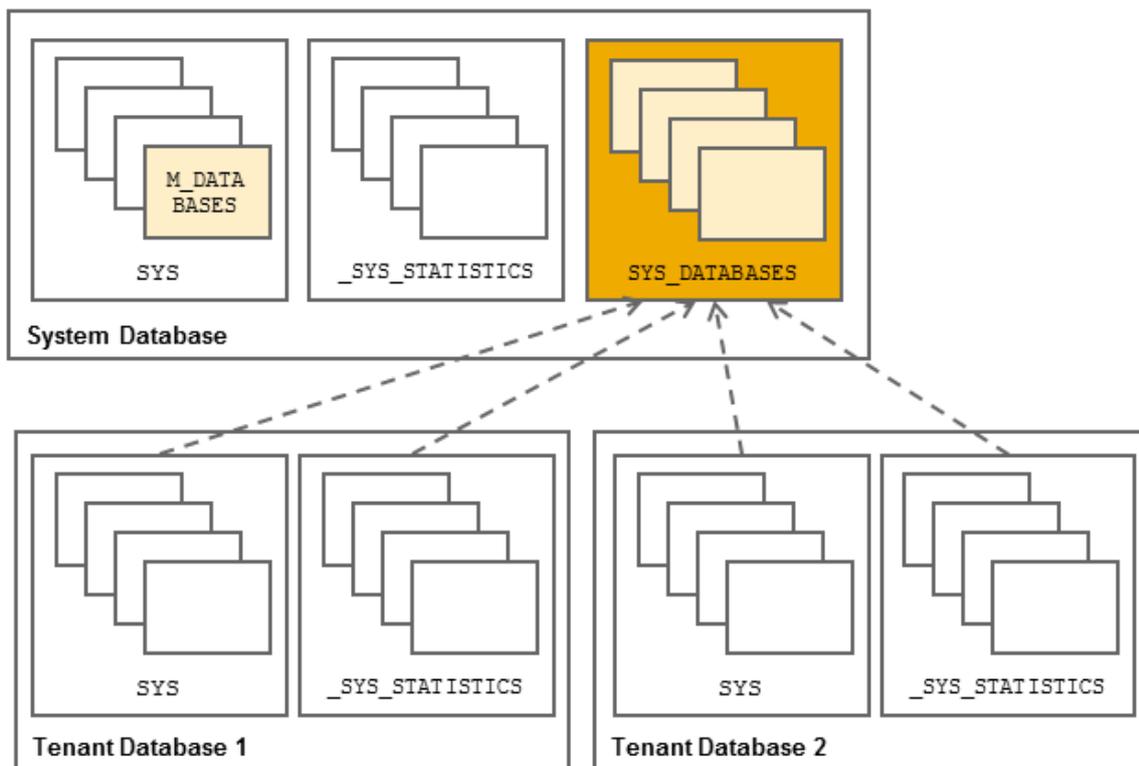
[Collecting Diagnosis Information for SAP Support \[page 484\]](#)

### 4.2.2.8 System and Statistics Views in Multiple-Container Systems

Every multitenant database container system 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 of multiple-container systems.



System and Statistics Views in Multiple-Container Systems

## 4.2.3 Managing Resources in Multiple-Container Systems

Manage and control the memory and CPU usage of your multiple-container 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. For more information about additional options, see the sections referenced under *Related Information*.

- `[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;
```

### 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;
```

### 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

[Parameter Reference: Memory Consumption \[page 281\]](#)

[Controlling Parallel Execution of SQL Statements \[page 427\]](#)

[Controlling CPU Consumption \[page 424\]](#)

[Configuring SAP HANA System Properties \(INI Files\) \[page 212\]](#)

## 4.2.4 Copying and Moving Tenant Databases Between Systems

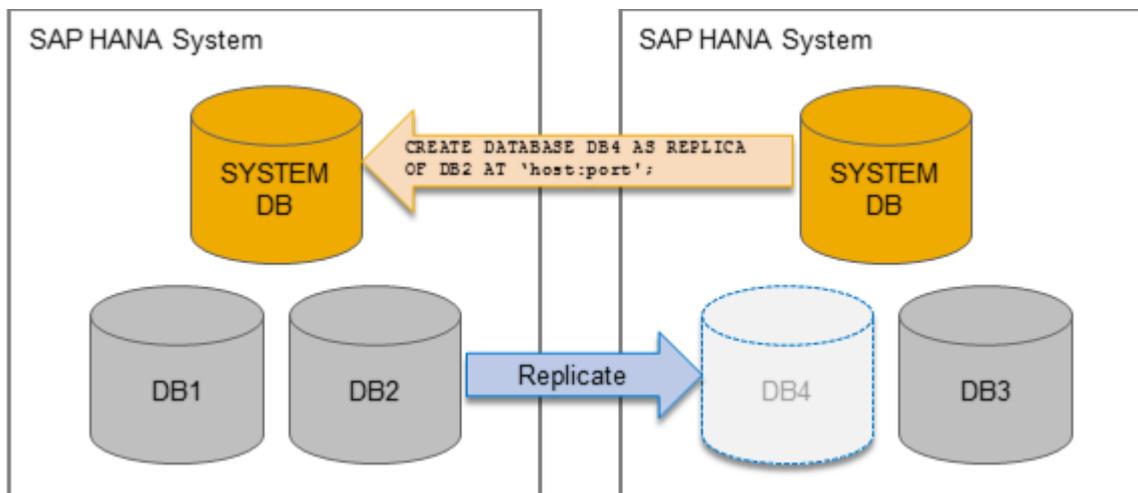
Using SAP HANA system replication mechanisms, SAP HANA multitenant database containers can be copied and moved securely and conveniently from one SAP HANA system to another with near-zero downtime. This allows you to respond flexibly to changing resource requirements and to manage your system landscape efficiently.

The following sections provide an overview of copying or moving a tenant database using system replication.

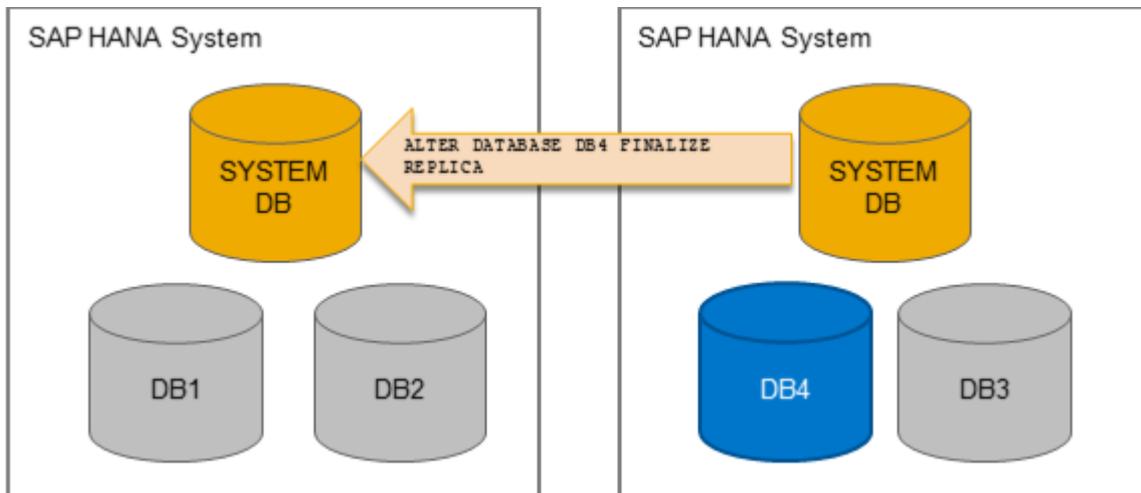
- [Process Overview \[page 145\]](#)
- [Use Cases \[page 146\]](#)
- [Which Data Is Copied or Moved? \[page 147\]](#)
- [Recoverability After Copy or Move \[page 148\]](#)
- [Prerequisites and Implementation Considerations \[page 148\]](#)
- [Other Copy and Move Methods \[page 148\]](#)

### Process Overview

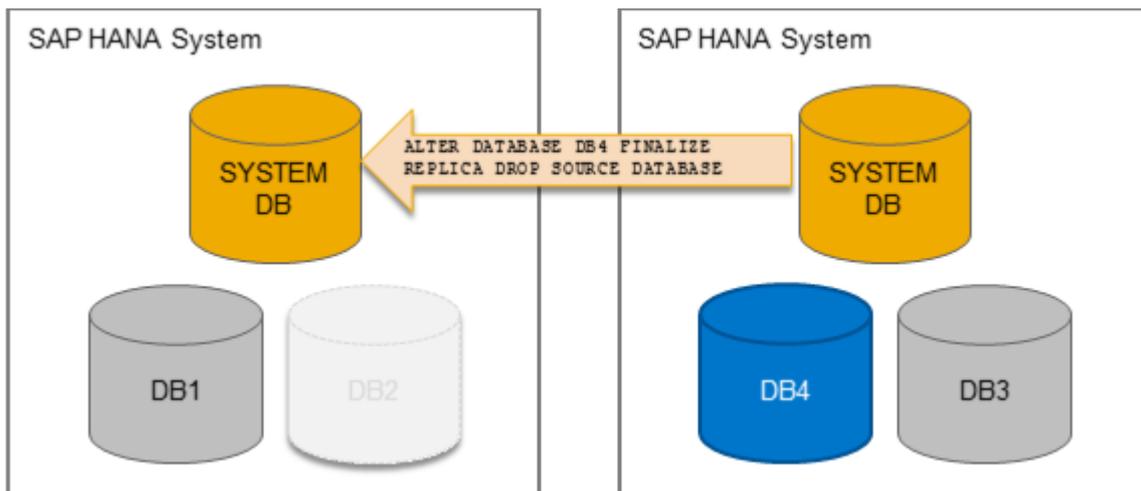
Copying and moving a tenant database are essentially the same process. First, the tenant database is copied through the replication of all of its data to a newly created tenant database in a target system:



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 to 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 This refers to files in the directory <code>\$DIR_INSTANCE/.. /SYS/global/hdb/custom/config/&lt;database_name&gt;</code>	No
Certificates and certificate collections stored in the tenant database 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.	Yes
<p><b>i Note</b></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><b>i Note</b></p> <p>The root key used for data volume encryption is not replicated.</p>	
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.

## 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.
- Since the copy and move process uses system replication mechanisms, system replication must not be enabled on either the source or target system for high availability purposes for the entire duration of the copy or move process.
- 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 smart data streaming (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 \[page 153\]](#)

[Copy and Move Process \[page 149\]](#)

[Copy or Clone an SAP HANA System \[page 594\]](#)

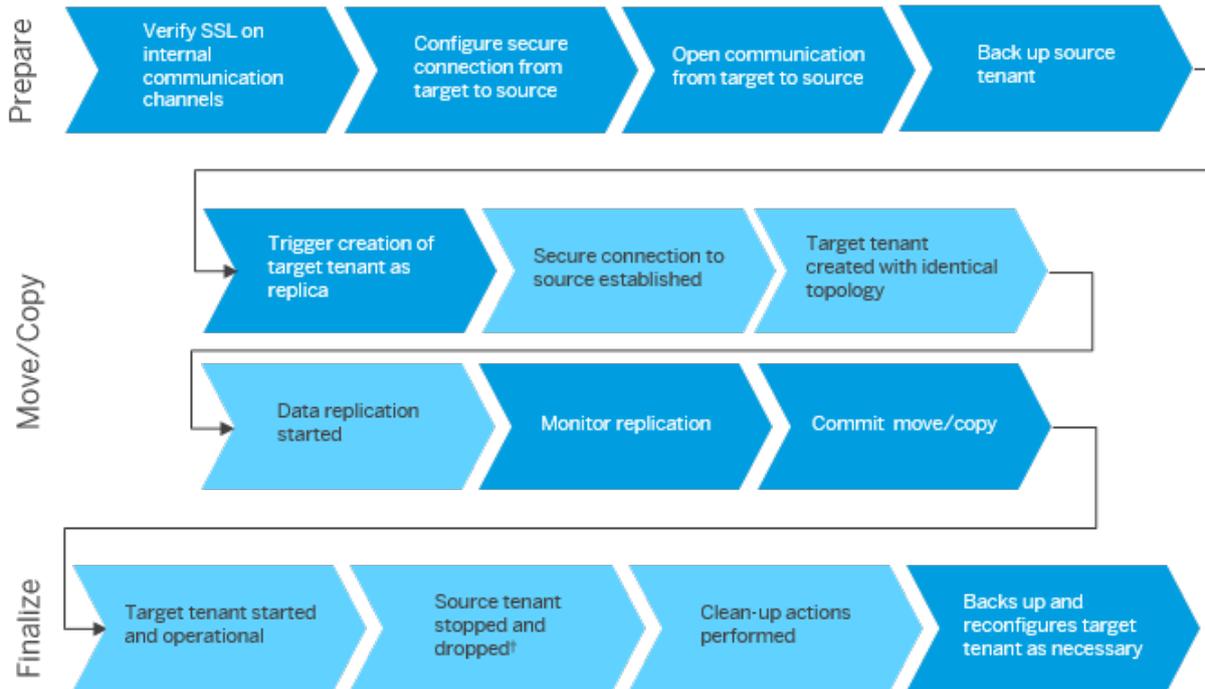
## **4.2.4.1 Copy and Move Process**

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
	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

Who?	Does What?	Where?
	<p>Configures secure connection from target system to source system by:</p> <ol style="list-style-type: none"> <li>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.</li> <li>2. Creating a credential in the target system to enable authenticated access from the target system to the source system</li> </ol>	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	<p>Establishes a secure connection to the system database using the stored credentials created above</p> <p>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.</p> <p>The generated key pairs and certificates are imported into two newly created certificate collections in the target system database.</p>	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?
	<p>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></p> <p>In the case of a move, the administrator indicates that the source tenant database is to be dropped: <code>DROP SOURCE DATABASE</code>.</p>	System database of the target system

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
	<p>Performs clean-up operations:</p> <ul style="list-style-type: none"> <li>Deletes any cross-database dependencies to the original tenant database in other tenant databases of the source system (<b>move only</b>)</li> <li>Deletes any remote identity dependencies of users in the new tenant database in the target tenant database (<b>copy and move</b>)</li> <li>Generates a new root key used for data volume encryption and re-encrypts data if data volume encryption is enabled (<b>copy and move</b>)</li> </ul>	Source system and tenant database in target system
System administrator	<p>Performs manual post-copy or post-move steps:</p> <ul style="list-style-type: none"> <li>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.</li> <li>Reverse preparatory steps required to secure the copy process</li> <li>If necessary, reconfigure parameters in *.ini files with tenant-database-specific values</li> <li>If necessary, reconfigure cross-database access</li> </ul>	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 &lt;database_name&gt;</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 \[page 153\]](#)

[Preparing to Copy or Move a Tenant Database \[page 155\]](#)

[Copy a Tenant Database to Another System \[page 162\]](#)

[Move a Tenant Database to Another System \[page 165\]](#)

### 4.2.4.2 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.

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 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.

### 4.2.4.3 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.

1. [Verify TLS/SSL Configuration of Internal Communication Channels \[page 155\]](#)  
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. [Open Communication From Target to Source System \[page 157\]](#)  
Open communication from the target system to the source system by enabling services in the source system to listen on all network interfaces.
3. [Set Up Trust Relationship Between Target and Source Systems \[page 158\]](#)  
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.
4. [Create Credential for Authenticated Access to Source System \[page 160\]](#)  
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 \[page 162\]](#)  
Back up the tenant database that will be copied or moved.

#### 4.2.4.3.1 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 \[page 155\]](#)

**Next task:** [Open Communication From Target to Source System \[page 157\]](#)

## 4.2.4.3.2 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 \[page 155\]](#)

**Previous task:** [Verify TLS/SSL Configuration of Internal Communication Channels \[page 155\]](#)

**Next task:** [Set Up Trust Relationship Between Target and Source Systems \[page 158\]](#)

### 4.2.4.3.3 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 -a <signature algorithm>:<key size>:<hash algorithm> -p  
<path>/<file name>.pse -x "" -noreq "CN=<source host name>"
```

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 \[page 155\]](#)

**Previous task:** [Open Communication From Target to Source System \[page 157\]](#)

**Next task:** [Create Credential for Authenticated Access to Source System \[page 160\]](#)

## Related Information

[SAP Note 2175664 - Migration of file system based X.509 certificate stores to in-database certificate stores](#)  
[Managing Client Certificates in the SAP HANA Database \[page 758\]](#)

### 4.2.4.3.4 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 \[page 155\]](#)

**Previous task:** [Set Up Trust Relationship Between Target and Source Systems \[page 158\]](#)

**Next task:** [Back Up Tenant Database \[page 162\]](#)

---

## 4.2.4.3.5 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 \[page 155\]](#)

**Previous task:** [Create Credential for Authenticated Access to Source System \[page 160\]](#)

### Related Information

[Creating Backups \[page 920\]](#)

## 4.2.4.4 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 (for example, in the SQL console of the SAP HANA studio):

### 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_s
ervice>' ] }...
{ AS REPLICAS 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.

### Sample Code

```
CREATE DATABASE TARGET_DATABASE AS REPLICAS 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 SYNC mode.
ERROR	A connection error occurred (details can be found in REPLICATION_STATUS_DETAILS).

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 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 \[page 145\]](#)

[Preparing to Copy or Move a Tenant Database \[page 155\]](#)

[Perform Manual Post-Copy/Move Tasks \[page 168\]](#)

### 4.2.4.5 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 (for example, in the SQL console of the SAP HANA studio):

#### 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_s
ervice>' ] }...
{ 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.

### 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 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 \[page 145\]](#)

[Preparing to Copy or Move a Tenant Database \[page 155\]](#)

[Perform Manual Post-Copy/Move Tasks \[page 168\]](#)

## 4.2.4.6 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 Multiple-Container Systems*.
5. Reconfigure cross-database access, if required. See *Enable and Configure Cross-Database Access*.

### Related Information

[Open Communication From Target to Source System \[page 157\]](#)

[Set Up Trust Relationship Between Target and Source Systems \[page 158\]](#)

[Create Credential for Authenticated Access to Source System \[page 160\]](#)

[Enable and Configure Cross-Database Access \[page 115\]](#)

[Configuration Parameters in Multiple-Container Systems \[page 213\]](#)

---

## 4.2.4.7 Tutorial: Moving a Tenant Database

In this tutorial, we'll move an existing tenant database from a test system to a production system.

### Scenario

M01 is a test system. A newly developed application is running here in the tenant database NEW\_APP\_TEST. After several rounds of successful application testing, it is time to move the application to the production system, M02. The production system already has some applications running in other tenant databases. Since we would like to have full database isolation between applications, the new application should also run in a separate tenant database in the production system. We will therefore move the tenant database NEW\_APP\_TEST in system M01 to system M02, where it will be called NEW\_APP\_PROD.

### Tools

To do the move, we will use several tools:

- Command line tool
- SAP HANA studio
- SAP Web Dispatcher Administration
- SAP HANA cockpit

### Authorization

We will need the following authorizations in both systems:

- Credentials of operating system users m01adm and m02adm
- Database user in the system database with the following system privileges and roles:
  - CATALOG READ
  - INIFILE ADMIN
  - SSL ADMIN
  - DATABASE ADMIN
  - BACKUP ADMIN
  - `sap.hana.xs.wdisp.admin::WebDispatcherAdmin`
  - `sap.hana.security.cockpit.roles::EditCertificateStore`
  - `sap.hana.security.cockpit.roles::MaintainCertificateCollections`
  - `sap.hana.security.cockpit.roles::MaintainCertificates`

As the target system is running in high isolation mode, you will also need root user access to be able to create an operating group and user for the new tenant database.

## Steps

This tutorial takes you through the following steps:

1. [Verify TLS/SSL Configuration of Internal Communication Channels in M01 and M02 \[page 170\]](#)  
Verify that TLS/SSL is enabled on internal communication channels in source system M01 and target system M02. M02 is running in high isolation mode, so this will automatically be the case.
2. [Open Communication From M02 to M01 \[page 171\]](#)  
Open communication from the target system M02 to the source system M01 by enabling services in the source system to listen on all network interfaces.
3. [Configure TLS/SSL for External Communication in M01 and M02 \[page 173\]](#)  
Configure the external communication channel for TLS/SSL in the source system M01 and the target system M02 using in-database certificate collections.
4. [Set Up Trust Relationship Between M02 and M01 \[page 176\]](#)  
Create an in-database certificate collection in the system database of the target system M02 and add the public-key certificate of the source system M01. During the move, this certificate will later be used to secure communication between the systems via external SQL connections.
5. [Create Credential for Authenticated Access from M02 to M01 \[page 179\]](#)  
Create a credential to enable authenticated access to the source system M01 from target system M02 for the purpose of moving the tenant database.
6. [Back Up Tenant Database NEW\\_APP\\_TEST \[page 180\]](#)  
Back up the tenant database NEW\_APP\_TEST in the source system.
7. [Create OS Group and OS for New Tenant Database in M02 \[page 181\]](#)  
Create an operating system group and user for the new tenant database that you will create in target system M02. This is necessary because M02 is configured for high isolation.
8. [Move Tenant Database NEW\\_APP\\_TEST \[page 181\]](#)  
Move the tenant database NEW\_APP\_TEST in the source system M01 to the target system M02 as the database NEW\_APP\_PROD. After the move, NEW\_APP\_TEST is deleted and NEW\_APP\_PROD is up and running.

### 4.2.4.7.1 Verify TLS/SSL Configuration of Internal Communication Channels in M01 and M02

Verify that TLS/SSL is enabled on internal communication channels in source system M01 and target system M02. M02 is running in high isolation mode, so this will automatically be the case.

## Procedure

In the system database of the source system (M01):

1. In the SAP HANA studio, open the *Configuration* tab of the Administration editor.
2. Navigate to the [communication] section of the `global.ini` file.

- Verify that value of the parameter `ssl` is set to **systemPKI** at the `SYSTEM` layer.

Name	Default	System
global.ini		◆
[ ] communication		◆
ssl	off	● systemPKI

### **i** Note

If **systemPKI** is not set, set it and restart the system.

In the system database of the target system (M02):

- Navigate to the `[communication]` section of the `global.ini` file.
- In the SAP HANA studio, open the *Configuration* tab of the Administration editor.
- Verify that value of the parameter `database_isolation` is set to **high** at the `SYSTEM` layer.

Name	Default	System
global.ini		◆
[ ] multidb		◆
database_isolation	low	● high

**Task overview:** [Tutorial: Moving a Tenant Database \[page 169\]](#)

**Next task:** [Open Communication From M02 to M01 \[page 171\]](#)

## 4.2.4.7.2 Open Communication From M02 to M01

Open communication from the target system M02 to the source system M01 by enabling services in the source system to listen on all network interfaces.

### Procedure

In the system database of the source system (M01):

- Open the SAP HANA database lifecycle manager by entering the following URL in a browser:  
`https://<host>:1129/lmsl/HDBLCM/M01/index.html`
- Click the tile *Configure Inter-Service Communication*.



3. When prompted, enter the password of the m01adm user.
4. Select the setting *global*.

**i Note**

You do not need to change any other settings.

5. Click *Run* to apply the new setting.
6. Close the application and log out.

In the system database of the target system (M02):

7. Repeat the steps above.

**Task overview:** [Tutorial: Moving a Tenant Database \[page 169\]](#)

**Previous task:** [Verify TLS/SSL Configuration of Internal Communication Channels in M01 and M02 \[page 170\]](#)

**Next task:** [Configure TLS/SSL for External Communication in M01 and M02 \[page 173\]](#)

## 4.2.4.7.3 Configure TLS/SSL for External Communication in M01 and M02

Configure the external communication channel for TLS/SSL in the source system M01 and the target system M02 using in-database certificate collections.

### Procedure

In the system database of the source system (M01):

1. Generate a new public and private key pair for the system.
  - a. Open the SAP Web Dispatcher Administration tool by entering the following URL in your browser:  
`http://<host_FQDN>:80<instance>/sap/hana/xs/wdisp/admin`
  - b. Navigate to the *PSE Management* screen and select the PSE `sapsrv.pse`.
  - c. Choose *Recreate PSE*.

The screenshot shows the 'Manage PSE' interface for 'sapsrv.pse'. At the top right, the 'Recreate PSE' button is highlighted with a green box. Below the header, there are buttons for 'Export Own Certificate', 'Create CA Request', and 'Import CA Response'. The 'PSE Attributes' section shows details like Subject, Issuer, Serialno, KeyInfo, Validity, KeyUsage, ExtKeyUsage, and SubjectAltName. The 'Trusted Certificates' section shows 'PKList is empty.' and an 'Import Certificate' button.

- d. Verify the information and choose *Create*.

The screenshot shows the 'Create new key pair for PSE sapsrv.pse' dialog box. The 'Create' button is highlighted with a green box. The dialog contains the following fields:

- Distinguished Name: `CN=..., OU=..., OU=HANA SSL`
- Algorithm: `RSA with SHA-256` (with an 'Algorithm Info' link)
- Key Length: `2048`
- PIN (optional): [Empty text box]

2. Extract the system's own certificate, certificate chain and private key from the PSE file `sapsrv.pse`:
  - a. Download the script `extract_certificates.py` from SAP Note [2175664](#) to the server.
  - b. From the command line, log on to the server as `m01adm` and navigate to the directory containing the PSE file `sapsrv.pse`:

```
cd /usr/sap/M01/HDB<instance>/mdchost/sec
```

- c. Execute the script as follows:

```
python <path_to_script> -p sapsrv.pse
```

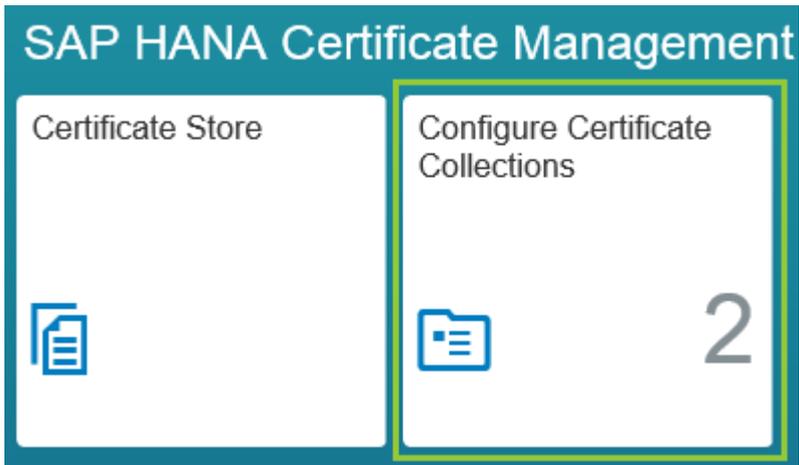
- d. Copy the output starting `-----BEGIN RSA PRIVATE KEY-----` and ending `-----END CERTIFICATE-----` to a text file.

```
:/usr/sap/ sec> python extract_certificates.py -p sapsrv.pse
ALTER PSE <name> SET OWN CERTIFICATE
'-----BEGIN RSA PRIVATE KEY-----
[REDACTED]
-----END RSA PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
[REDACTED]
-----END CERTIFICATE-----';
```

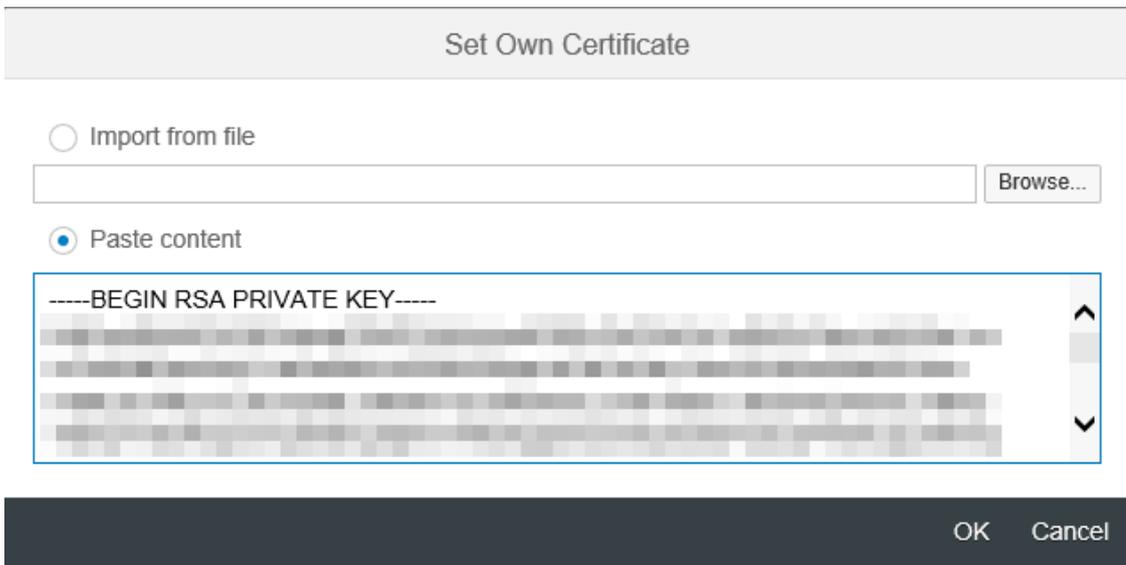
**⚠ Caution**

The extraction process prints the private key directly in PEM form to the screen. Make sure that no unauthorized access is possible to this key information because this could compromise the security of the own certificate and the private key.

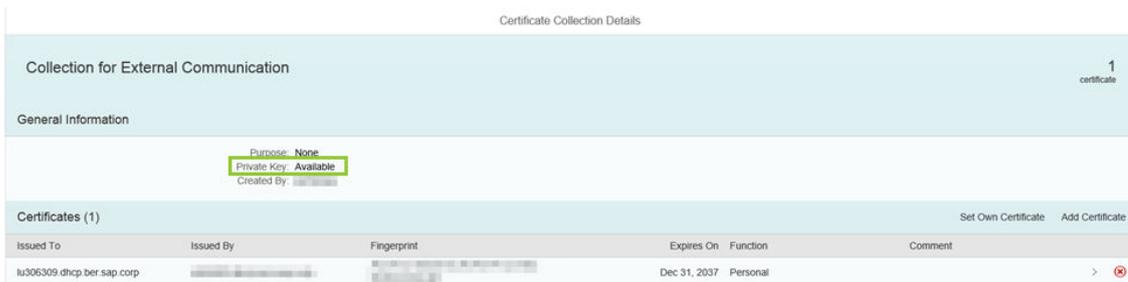
3. Create a certificate collection for external communication:
  - a. Open the SAP HANA cockpit by entering the following URL in your browser:  
`http://<host_FQDN>:80<instance>/sap/hana/admin/cockpit`
  - b. Open the *Configure Certificate Collections* app:



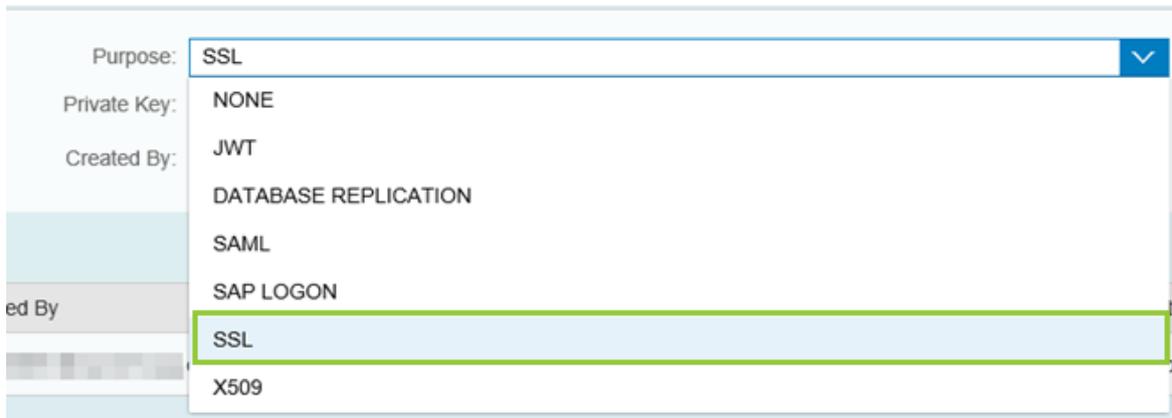
- c. Create a new collection.
- d. In the new collection, click *Set Own Certificate* and paste in the extracted certificate contents:



Notice that it is now indicated that the collection has a private key available:



- 4. Set the purpose of the collection to *SSL*:



5. Save the collection and exit the application.

In the system database of the target system:

6. Repeat the steps above.

**Task overview:** [Tutorial: Moving a Tenant Database \[page 169\]](#)

**Previous task:** [Open Communication From M02 to M01 \[page 171\]](#)

**Next task:** [Set Up Trust Relationship Between M02 and M01 \[page 176\]](#)

## 4.2.4.7.4 Set Up Trust Relationship Between M02 and M01

Create an in-database certificate collection in the system database of the target system M02 and add the public-key certificate of the source system M01. During the move, this certificate will later be used to secure communication between the systems via external SQL connections.

### Procedure

In the system database of the source system (M01):

1. Get the public-key certificate of the source system.
  - a. Open the SAP Web Dispatcher Administration tool by entering the following URL in your browser:  
`http://<host_FQDN>:80<instance>/sap/hana/xs/wdisp/admin`
  - b. Navigate to the *PSE Management* screen and select the PSE `sapsrv.pse`.
  - c. Click *Export Own Certificate*.

**Manage PSE** sapsrv.pse PSE Info [Recreate PSE](#) [Delete PSE](#) | [Create New PSE](#)

Version of Security Library (SSL/TLS): CommonCryptoLib Version 8.4.49 Mar 4 2016  
 Location of Personal Security Environment (PSE) Files: /usr/sap/.../sec

**PSE Attributes** [Export Own Certificate](#) [Create CA Request](#) [Import CA Response](#)

```

Subject      : CN=..., OU=MD1, OU=HANA SSL
Issuer       : CN=..., OU=MD1, OU=HANA SSL
Serialno     : OA:
KeyInfo      : RSA, 2048-bit
Validity -   NotBefore: Mon Jun 13 12:48:52 2016 (160613114852Z)
              NotAfter : Fri Jan 1 01:00:01 2038 (380101000001Z)
KeyUsage     : none
ExtKeyUsage  : none
SubjectAltName : none
  
```

**Trusted Certificates** [Import Certificate](#)

PKList is empty.

d. Copy the certificate contents to a text file.

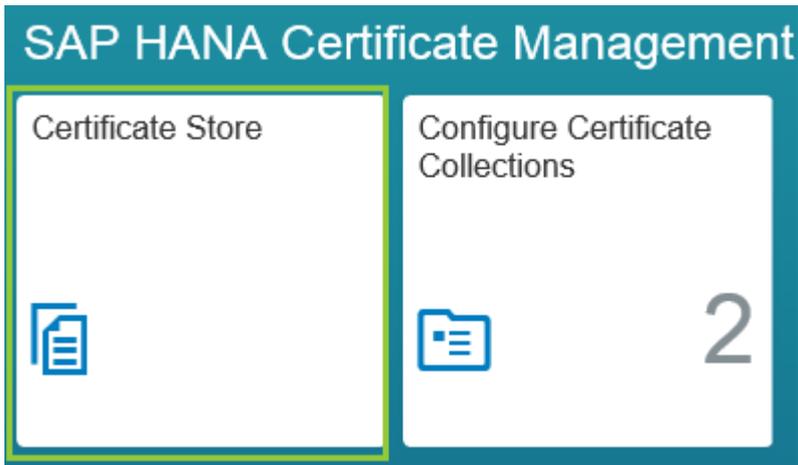
**Certificate of PSE SAPSSLS.pse** [Back](#)

```

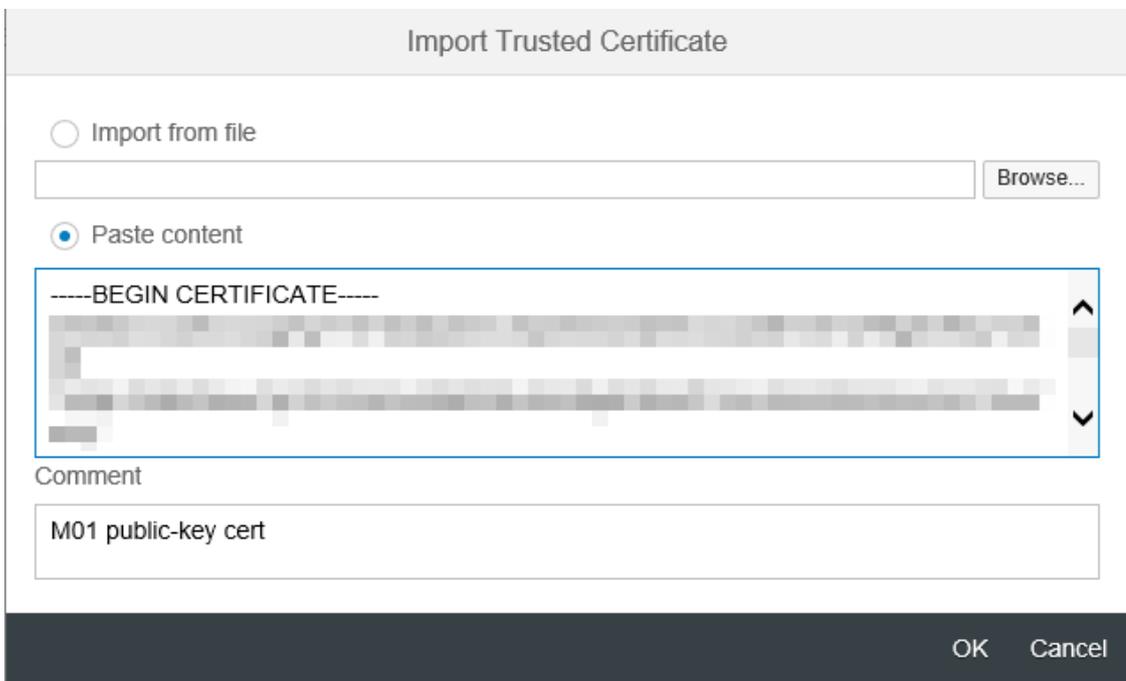
-----BEGIN CERTIFICATE-----
[Redacted Certificate Content]
-----END CERTIFICATE-----
  
```

In the system database of the target system (M02):

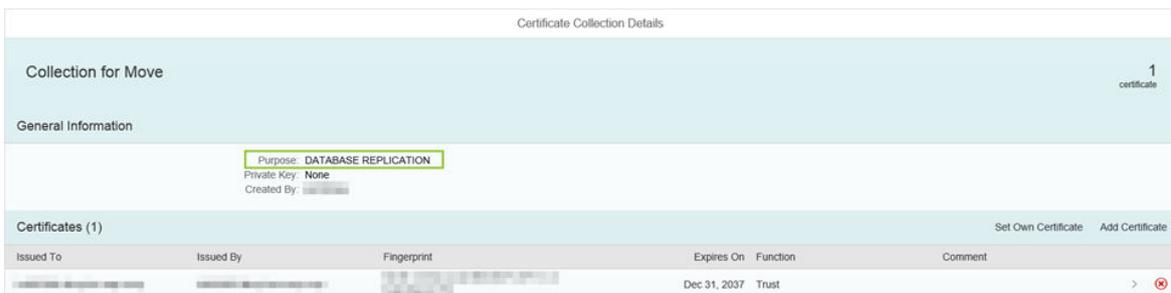
2. Import the public-key certificate of the source system into the certificate store:
  - a. Open the SAP HANA cockpit by entering the following URL in your browser:  
`http://<host_FQDN>:80<instance>/sap/hana/admin/cockpit`
  - b. Open the *Certificate Store* app.



- c. Click *Import* and paste the contents of the source system certificate.



3. Create a certificate collection for the purpose of database replication.
  - a. Open the *Configure Certificate Collections* app.
  - b. Create a new collection.
  - c. Click *Add Certificate* and select the public-key certificate of the source system.
4. Set the purpose of the collection to *DATABASE REPLICATION* and save.



---

**Task overview:** [Tutorial: Moving a Tenant Database \[page 169\]](#)

**Previous task:** [Configure TLS/SSL for External Communication in M01 and M02 \[page 173\]](#)

**Next task:** [Create Credential for Authenticated Access from M02 to M01 \[page 179\]](#)

## 4.2.4.7.5 Create Credential for Authenticated Access from M02 to M01

Create a credential to enable authenticated access to the source system M01 from target system M02 for the purpose of moving the tenant database.

### Procedure

In the system database of the target system (M02):

1. In the SAP HANA studio, open the SQL console.
2. Execute the following statement:

```
CREATE CREDENTIAL FOR COMPONENT 'DATABASE_REPLICATION' PURPOSE
'<host>:<internal_port>' TYPE 'PASSWORD' USING
'user="<database_administrator_user>";password="<password>"'
```

#### **i** Note

As the PURPOSE you enter the host name and internal port number of the system database of the source system M01. The user to be specified after the keyword USING is a user in the system database of the source system M01 with the system privilege DATABASE ADMIN.

**Task overview:** [Tutorial: Moving a Tenant Database \[page 169\]](#)

**Previous task:** [Set Up Trust Relationship Between M02 and M01 \[page 176\]](#)

**Next task:** [Back Up Tenant Database NEW\\_APP\\_TEST \[page 180\]](#)

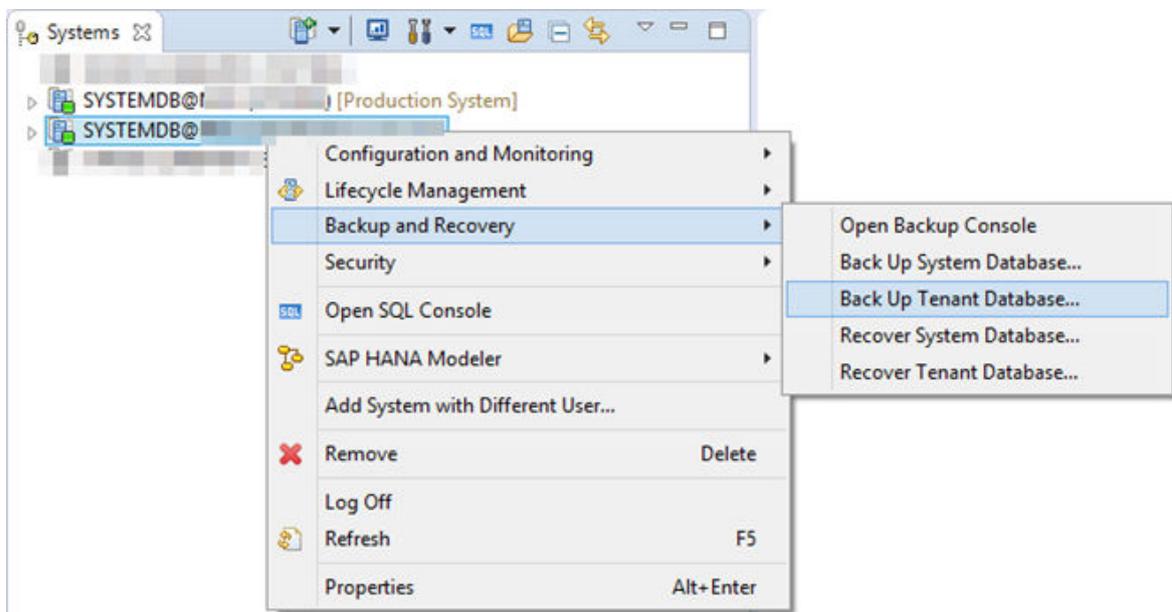
## 4.2.4.7.6 Back Up Tenant Database NEW\_APP\_TEST

Back up the tenant database NEW\_APP\_TEST in the source system.

### Procedure

In the system database of the source system (M01):

1. In the SAP HANA studio, right-click the source system and choose ► *Backup and Recovery* ► *Backup Tenant Database* ►.



2. Select the tenant database NEW\_APP\_TEST.
3. Select *Complete Data Backup*.
4. Specify the destination type and location for the backup.
5. Specify the backup prefix.
6. Choose *Next* and *Finish*.

**Task overview:** [Tutorial: Moving a Tenant Database \[page 169\]](#)

**Previous task:** [Create Credential for Authenticated Access from M02 to M01 \[page 179\]](#)

**Next task:** [Create OS Group and OS for New Tenant Database in M02 \[page 181\]](#)

## 4.2.4.7.7 Create OS Group and OS for New Tenant Database in M02

Create an operating system group and user for the new tenant database that you will create in target system M02. This is necessary because M02 is configured for high isolation.

### Procedure

In the target system (M02):

1. As root user, log on to the server on which the name server of the system database is running.
2. Create a new group:

```
groupadd new_app_prod
```

3. Create a new user, specifying `sapsys` as the primary group:

```
useradd -g sapsys new_app_prod
```

4. Add the new user to the `m02shm` group and their own group as secondary groups:

```
usermod -G m02shm,new_app_prod new_app_prod
```

**Task overview:** [Tutorial: Moving a Tenant Database \[page 169\]](#)

**Previous task:** [Back Up Tenant Database NEW\\_APP\\_TEST \[page 180\]](#)

**Next task:** [Move Tenant Database NEW\\_APP\\_TEST \[page 181\]](#)

## 4.2.4.7.8 Move Tenant Database NEW\_APP\_TEST

Move the tenant database `NEW_APP_TEST` in the source system M01 to the target system M02 as the database `NEW_APP_PROD`. After the move, `NEW_APP_TEST` is deleted and `NEW_APP_PROD` is up and running.

### Procedure

In the system database of the target system (M02):

1. In the SAP HANA studio, open the SQL console.

2. Create the new tenant database NEW\_APP\_PROD in the target system as a copy of the original tenant database NEW\_APP\_TEST in the source system.

```
CREATE DATABASE NEW_APP_PROD AS REPLICA OF NEW_APP_TEST AT
'<host_name>:<port_number>' OS USER 'new_app_prod' OS GROUP 'new_app_prod'
```

3. Monitor replication progress of data replication from the NEW\_APP\_TEST to NEW\_APP\_PROD.

```
SELECT * FROM "SYS"."M_DATABASE_REPLICAS"
```

4. When replication status is ACTIVE, stop replication and finalize the move by executing the following statement:

```
ALTER DATABASE NEW_APP_PROD FINALIZE REPLICA DROP SOURCE DATABASE
```

**Task overview:** [Tutorial: Moving a Tenant Database \[page 169\]](#)

**Previous task:** [Create OS Group and OS for New Tenant Database in M02 \[page 181\]](#)

## 4.2.5 Using SAP Web Dispatcher for Load Balancing with Tenant Databases

If an SAP HANA system has multiple instances of SAP HANA extended services (SAP HANA XS) 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. This is also possible for systems with tenant databases, but requires additional configuration.

SAP Note [1855097](#) describes how to configure an external SAP Web Dispatcher for single-container SAP HANA systems. The following sections describe the additional configuration required for 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. For use with multiple-container systems, 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 Multitenant Database Containers*.
- 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 Multitenant Database Containers* (internal Web Dispatcher configuration) and *Configuring SAP Web Dispatcher to Support SSL* in the SAP HANA Web Dispatcher documentation.

## Related Information

[SAP Note 1855097](#)

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

[Configuring SAP Web Dispatcher to Support SSL](#)

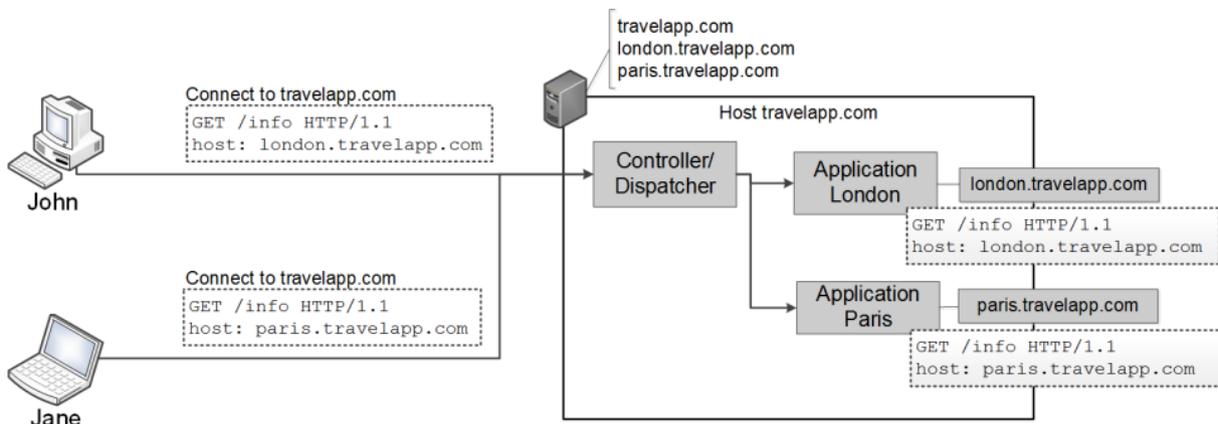
[Virtual-Host-Based Routing \[page 183\]](#)

[Configuring an External SAP Web Dispatcher for Tenant Databases \[page 185\]](#)

### 4.2.5.1 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:



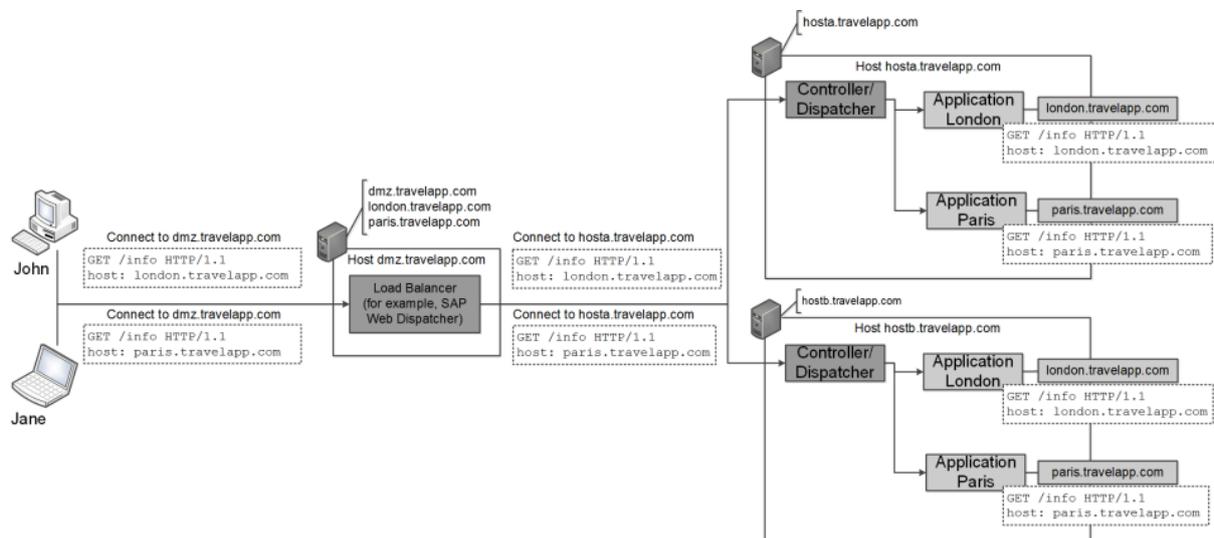
Virtual-Host-Based Routing

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:

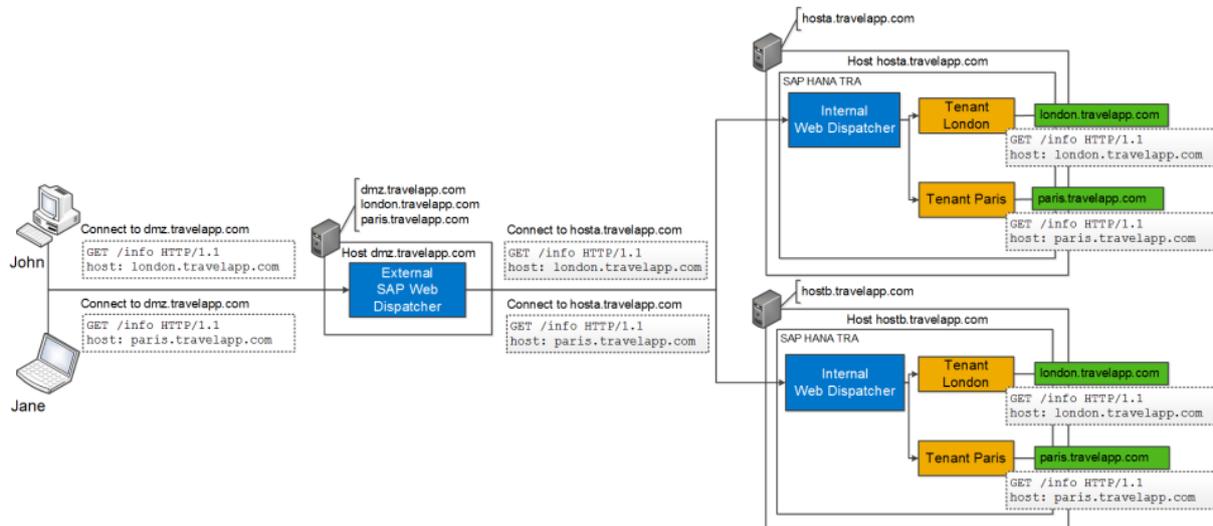


Virtual-Host-Based Routing with Load Balancing

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 Multitenant Database Containers

Translated to the context of SAP HANA multitenant database containers, 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 Multitenant Database Containers

### 4.2.5.2 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 Multitenant Database Containers*.

### Single Versus Multiple Tenant Access via External Web Dispatcher

Every tenant database that needs to be accessed through the external Web Dispatcher (and your system is running SPS 11 or higher) 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.

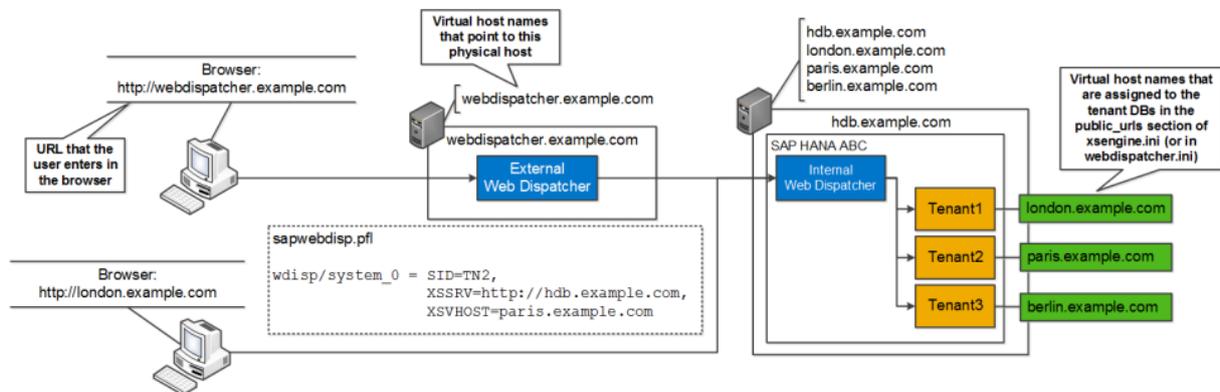
## Example

```
wdisp/system_<xx> = SID=<3-digit ID>, XSSRV=http://<physical host name of SAP HANA server>:<port>, XSVHOST=<virtual host name of the tenant>
```

## 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 Multitenant Database Containers*.

If only **one tenant database** needs to be accessed through the external Web Dispatcher (and your system is running SPS 11 or higher), a single `wdisp/system_<XX>` entry for the tenant database with the above configuration is sufficient, as depicted in the following figure:



## 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 (SPS 11 and higher)

## 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.

If **all or multiple tenant databases** need to be accessed (and your system is running SPS 09 or SPS 10), in addition to a `wdisp/system_<XX>` parameter entry for each tenant database, it is necessary to configure the external Web Dispatcher to differentiate between the various tenant databases. Since each `wdisp/system_<xx>` entry in the external Web Dispatcher represents one tenant database, incoming HTTP requests have to be mapped to a tenant database, that is to the right `wdisp/system_<xx>` entry in the internal Web Dispatcher configuration.

Virtual hosts 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 Multitenant Database Containers \[page 1070\]](#)

[Option 1: Configuring Access to Multiple \(or All\) Tenant Databases Through External Web Dispatcher Only \[page 187\]](#)

[Option 2: Configuring Access to Multiple \(or All\) Tenant Databases Through External Web Dispatcher and Directly \[page 190\]](#)

### 4.2.5.2.1 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 Multitenant Database Containers*.

## 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

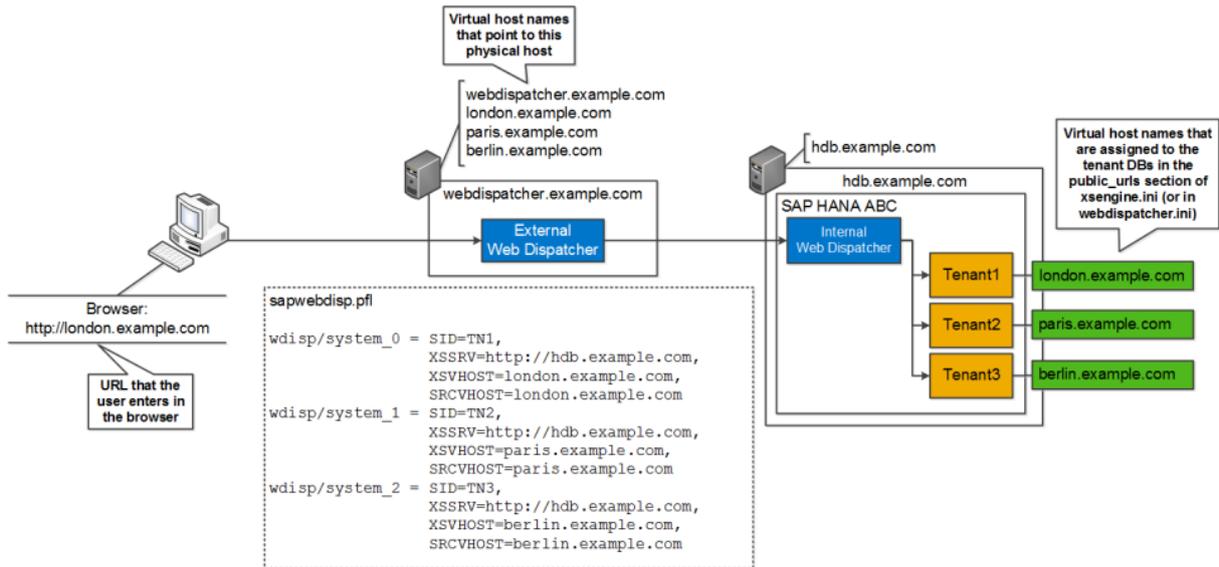
If a tenant database has multiple virtual host names assigned, only one needs to be entered in `XSVHOST`.

- SRCVHOST specifies the virtual host name that is used to map incoming HTTP requests to the `wdisp/system` entry that represents a particular tenant.

### 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 `XSRV` 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 `XSRV` 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.

### **i** Note

The injected HTTP header field can only be used in systems running SPS 11 or higher. Because of this, additional manual configuration is required for SPS 9 and SPS 10. See *Additional Configuration Required in SPS 09 and SPS 10*.

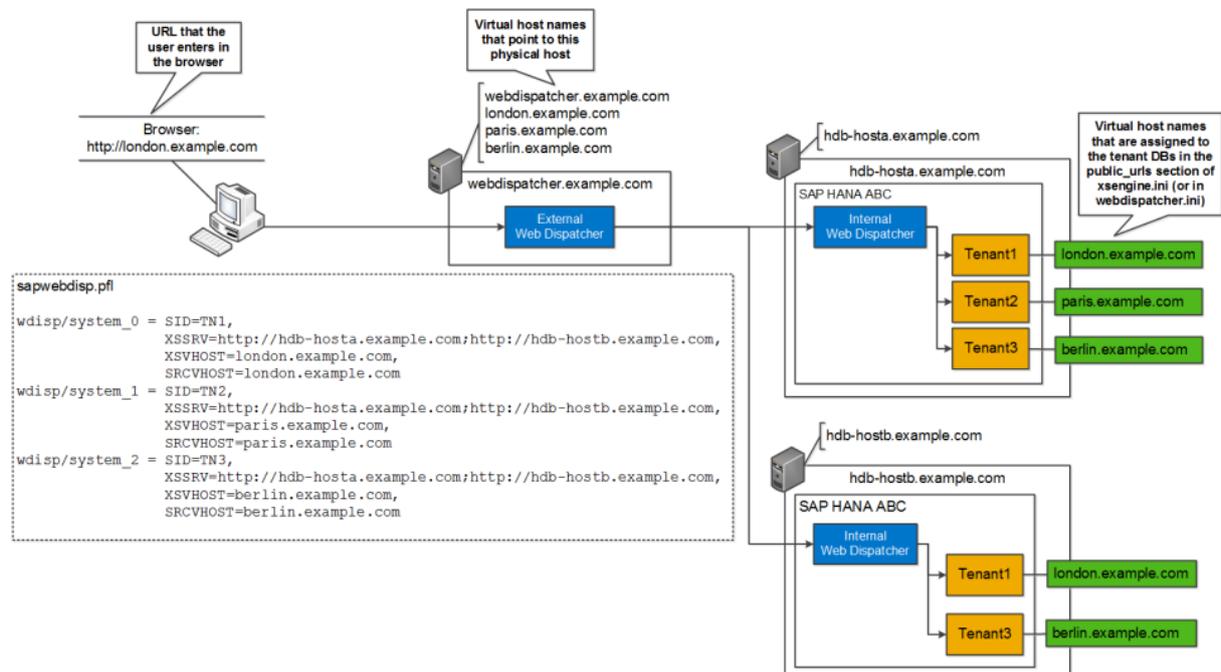
## Multiple-Host Systems

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

[Additional Configuration Required in SPS 09 and SPS 10 \[page 193\]](#)

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

### 4.2.5.2.2 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 Multitenant Database Containers*.

## Single-Host Systems

The `wdisp/system_<xx>` entry for each tenant database is then configured in the external Web Dispatcher profile as follows:

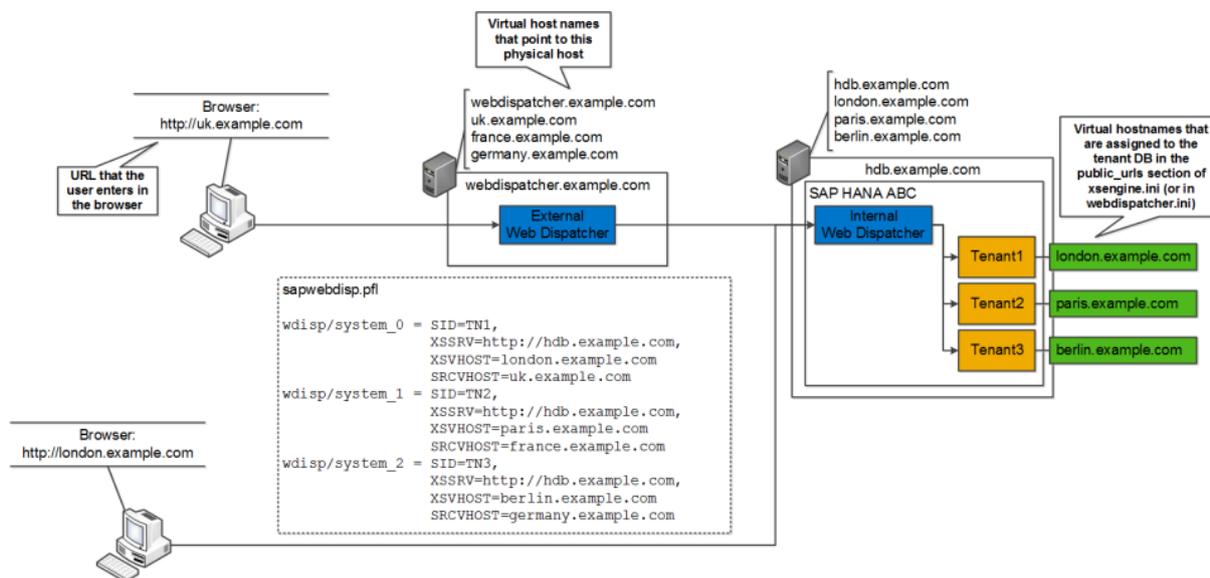
- `XSSRV` specifies the actual physical host name and port to which requests are sent.
- `XSVHOST` specifies the virtual host name of the tenant database to which requests are sent.

#### **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 (SPS 11 and Higher)

### 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.

### Note

The injected HTTP header field can only be used in SPS 11 or higher. Because of this additional manual configuration is required for SPS 9 and SPS 10. See *Additional Configuration Required in SPS 09 and SPS 10*.

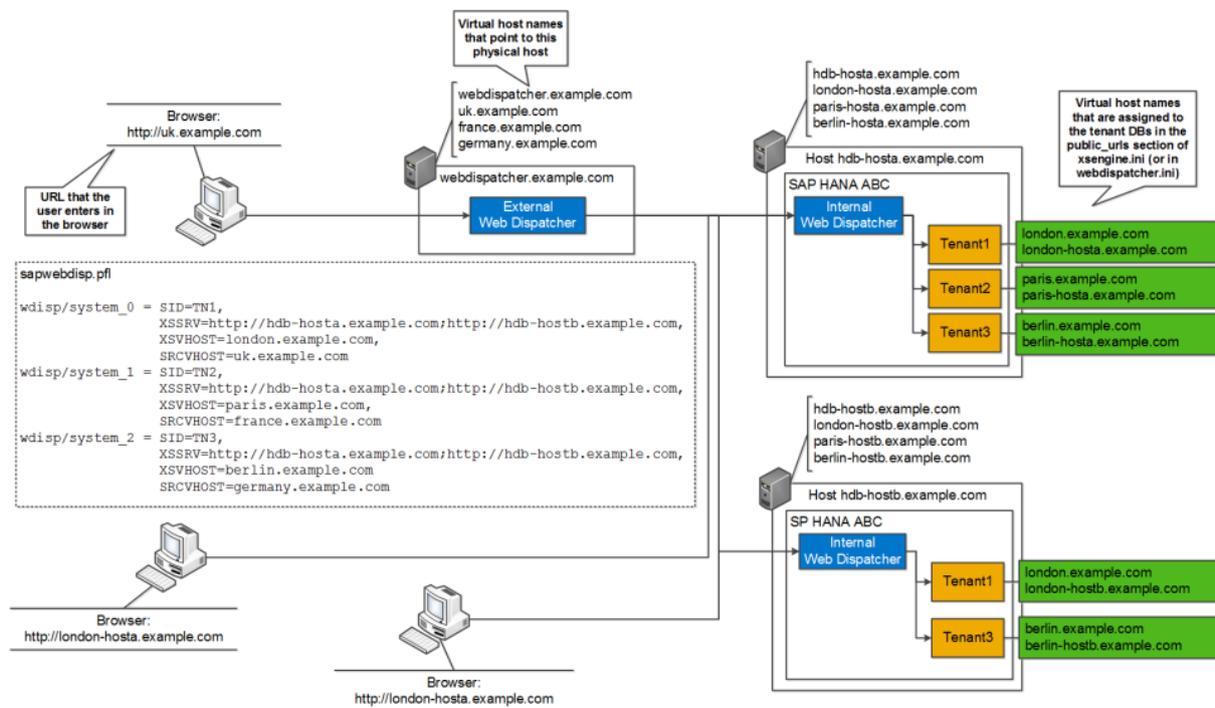
## Multiple-Host Systems

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 (SPS 11 and Higher)

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

[Additional Configuration Required in SPS 09 and SPS 10 \[page 193\]](#)

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

### 4.2.5.2.3 Additional Configuration Required in SPS 09 and SPS 10

In SAP HANA systems running SPS 09 and SPS 10, the configurations described in the previous sections are valid. However, an additional step is required for some access scenarios.

Additional configuration is required for the following access scenarios:

- Single tenant database access through an external Web Dispatcher (access through the external Web Dispatcher and directly)
- Multiple tenant database access through an external SAP Web Dispatcher and directly (that is, configuration option 2 for multiple tenant database access)

Additional configuration is not required for access to multiple tenant databases through an external Web Dispatcher only (that is, configuration option 1 for multiple tenant database access).

#### Note

This additional configuration does not have to be reverted after you upgrade from SPS 09 or SPS 10 to SPS 11 or higher.

## Access to Single Tenant Database

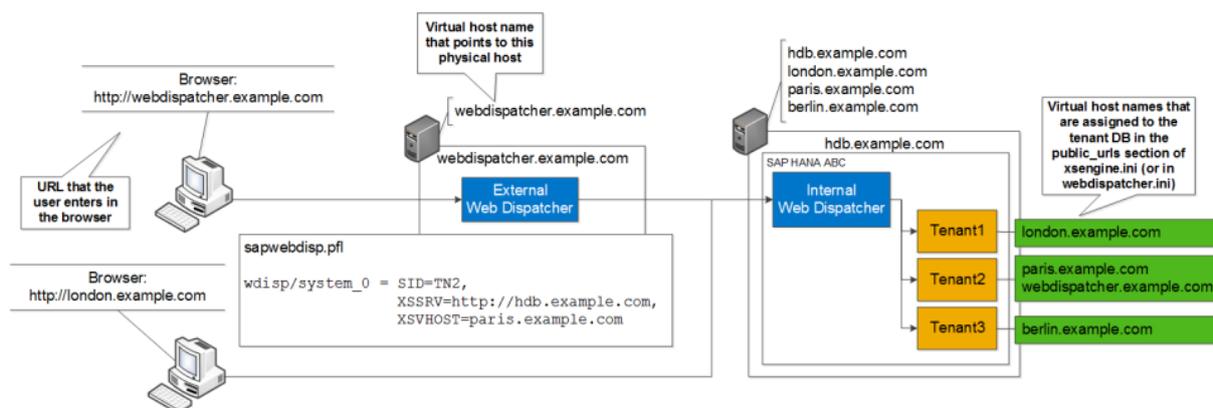
For access to a single tenant database, the host name of the external Web Dispatcher must be added to the SRCVHOST subparameter of the wdisp/system parameter for the tenant database in the webdispatcher.ini configuration file.

In SPS 09, you do this by editing the webdispatcher.ini file directly.

#### Example

```
ALTER SYSTEM ALTER CONFIGURATION ('webdispatcher.ini', 'system', '<tenant_DB_name>') SET (wdisp/system_xx) = 'SID=<3-character ID>, EXTSRV=http://localhost:3$(SAPSYSTEM)<internal xsengine port>, SRCVHOST=<Tenant DB FQDN>;<host name of external Web Dispatcher>' WITH RECONFIGURE;
```

In **SPS 10**, you add the host name to the public\_urls section of the xsengine.ini configuration file, for example, http\_webdispatcher\_url=http://<hostname>:80\$(SAPSYSTEM).



### 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. Also note that in SPS 09, the internal Web Dispatcher configuration must be done directly in the `webdispatcher.ini` file. Configuration using the `public_urls` section of the `xsengine.ini` file was introduced with SPS 10.

#### Access to a Single Tenant Database (SPS 09 and 10)

## Access to Multiple Tenant Databases

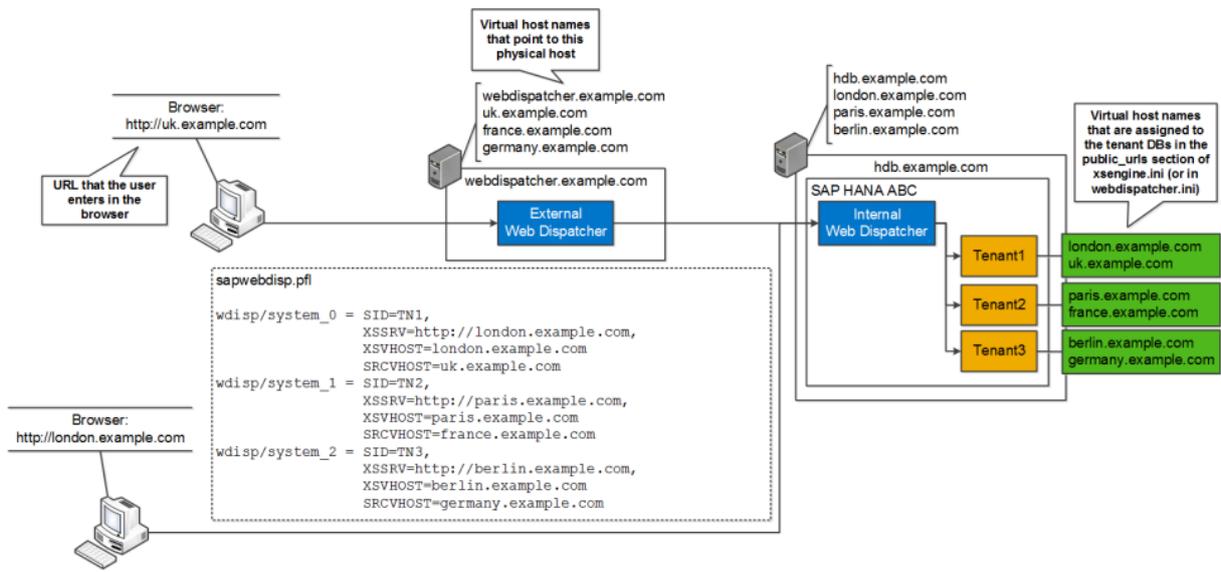
For access to multiple tenant databases, the additional virtual host names of the tenant databases configured in the external Web Dispatcher profile (that is the host name specified in the `SRCVHOST` subparameter of the `wdisp/system` parameter for a tenant database) must be added to the `SRCVHOST` subparameter of the `wdisp/system` parameter for each tenant databases in the `webdispatcher.ini` configuration file.

In SPS 09, you do this by editing the `webdispatcher.ini` file directly:

### Example

```
ALTER SYSTEM ALTER CONFIGURATION ('webdispatcher.ini', 'system', '<tenant_DB_name>') SET (wdisp/
system_xx) = 'SID=<3-character ID>, EXTSRV=http://localhost:3$(SAPSYSTEM)<internal
xsengine port>, SRCVHOST=<Tenant DB FQDN>;<additional virtual host name>' WITH
RECONFIGURE;
```

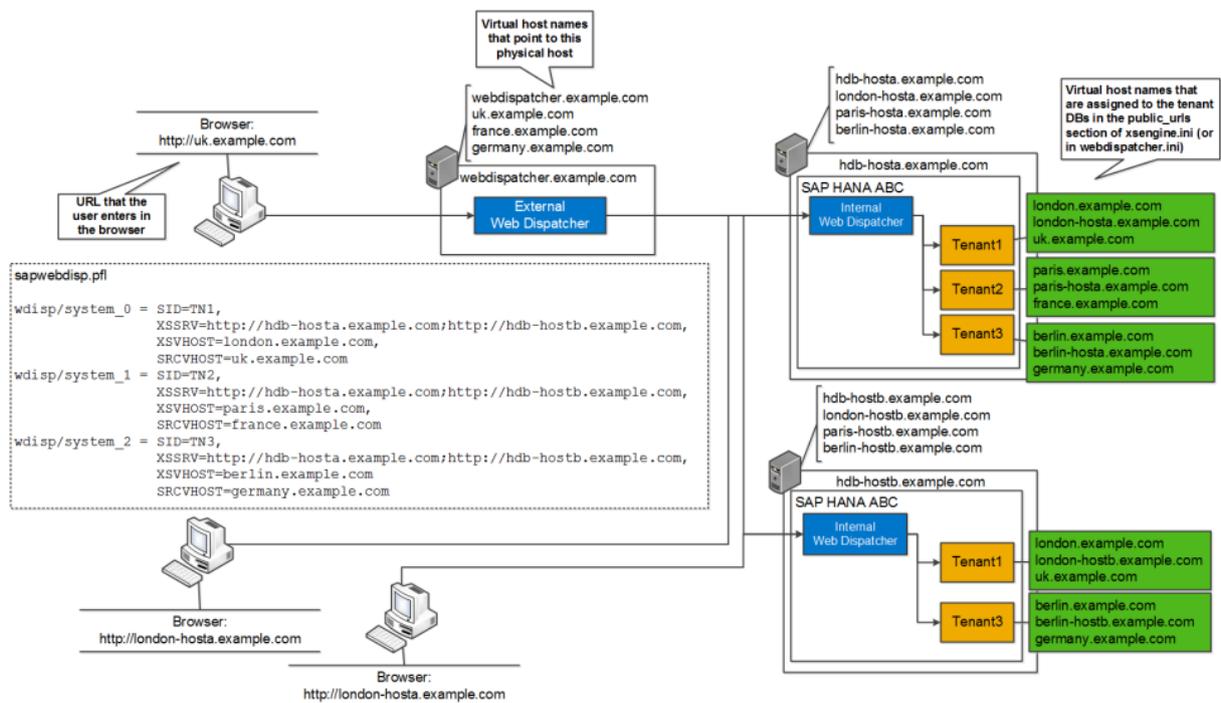
In **SPS 10**, you add the host name to the `public_urls` section of the `xsengine.ini` configuration file, for example, `http_webdispatcher_url=http://<hostname>:80$(SAPSYSTEM)`.



### 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. Also note that in SPS 09, the internal Web Dispatcher configuration must be done directly in the `webdispatcher.ini` file. Configuration using the `public_urls` section of the `xsengine.ini` file was introduced with SPS 10.

### Access to Multiple Tenant Databases in Single-Host System (SPS 09 and 10)



### **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. Also note that in SPS 09, the internal Web Dispatcher configuration must be done directly in the `webdispatcher.ini` file. Configuration using the `public_urls` section of the `xsengine.ini` file was introduced with SPS 10.

Access to Multiple Tenant Databases in Multiple-Host System (SPS 09 and 10)

## 4.2.6 Tutorial: Migrating SAP DB Control Center to a Tenant Database

SAP DB Control Center (SAP DCC) is an application that runs on SAP HANA. In this tutorial, you'll migrate an SAP DCC installation so that it goes from running on a single-container SAP HANA system to running in the tenant database of the same system converted to support multitenancy.

### Context

This tutorial takes you through the following steps:

1. [Install an SAP HANA System in Single-Container Mode \[page 197\]](#)  
In this step, you'll install an SAP HANA system, `MD0`, in single-container mode using the SAP HANA database lifecycle manager. This is the system on which you'll then install SAP DCC.
2. [Install and Configure SAP DCC on Single-Container System \[page 201\]](#)  
In this step, you'll install SAP DCC by importing the corresponding delivery unit (DU) into the newly installed system. Then, you'll configure SAP DCC and add systems to be monitored.
3. [Convert and Configure System for Multitenancy \[page 205\]](#)  
In this step, you'll convert the `MD0` system on which SAP DCC is running into a multiple-container system using the SAP HANA database lifecycle manager (HDBLCM). Afterward, SAP DCC will be running on the first tenant database in the multitenant system.

## 4.2.6.1 Install an SAP HANA System in Single-Container Mode

In this step, you'll install an SAP HANA system, MD0, in single-container mode using the SAP HANA database lifecycle manager. This is the system on which you'll then install SAP DCC.

### Prerequisites

- You're logged on as root user.
- You're familiar with installation tools and procedures as documented in the *SAP HANA Server Installation and Update Guide*.

### Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	<pre>cd &lt;installation medium&gt;/DATA_UNITS/ HDB_LCM_LINUX_X86_64</pre>
IBM Power Systems	<pre>cd &lt;installation medium&gt;/DATA_UNITS/ HDB_LCM_LINUX_PPC64</pre>

#### **i** Note

If you downloaded the components to a different directory, change to the directory where you unpacked the archive.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblcm
```

#### **i** Note

You can also use the graphical user interface tool `./hdblcgui`.

3. Select the index for *Install New System*, then select .
4. Select the components you would like to install as a comma-separated list, then select .
5. Specify the SAP HANA system properties when prompted.

In particular, select the *single\_container* value for the *Database Mode* property. For more information about the other system properties, see the *SAP HANA Server Installation and Update Guide*.

```

SAP HANA Lifecycle Management - Database Installation 1.00.100.00.1434444496
*****

Enter Local Host Name [sap-2020-00-00]:
Enter Installation Path [/hana/shared]:
Enter SAP HANA System ID: MD1
Enter Instance Number [00]: 01

-----
Index | Database Mode | Description
-----
1 | single_container | The system contains one database
2 | multiple_containers | The system contains one system database and 1..n tenant databases
-----
Select Database Mode / Enter Index [1]: 1

-----
Index | System Usage | Description
-----
1 | production | System is used in a production environment
2 | test | System is used for testing, not production
3 | development | System is used for development, not production
4 | custom | System usage is neither production, test nor development
-----
Select System Usage / Enter Index [4]: 3
Enter System Administrator (nd1adm) Password:
Confirm System Administrator (nd1adm) Password:
Enter System Administrator Home Directory [/usr/sap/MD1/home]:
Enter System Administrator Login Shell [/bin/sh]:
Enter System Administrator User ID [1000]:
Enter Location of Data Volumes [/hana/shared/MD1/global/hdb/data]:
Enter Location of Log Volumes [/hana/shared/MD1/global/hdb/log]:
Restrict maximum memory allocation? [n]:
Enter Database User (SYSTEM) Password:
Confirm Database User (SYSTEM) Password:
Restart instance after machine reboot? [n]:

Summary before execution:
Installation Path: /hana/shared
SAP HANA System ID: MD1
Instance Number: 01
Database Mode: single_container
System Usage: development
System Administrator Home Directory: /usr/sap/MD1/home
System Administrator Login Shell: /bin/sh
System Administrator User ID: 1000
ID of User Group (sapsys): 79
Location of Data Volumes: /hana/shared/MD1/global/hdb/data
Location of Log Volumes: /hana/shared/MD1/global/hdb/log
Local Host Name: sap-2020-00-00

Do you want to continue? (y/n): y

```

- Once you've provided the installer with all requested responses, enter **y** to continue:

```

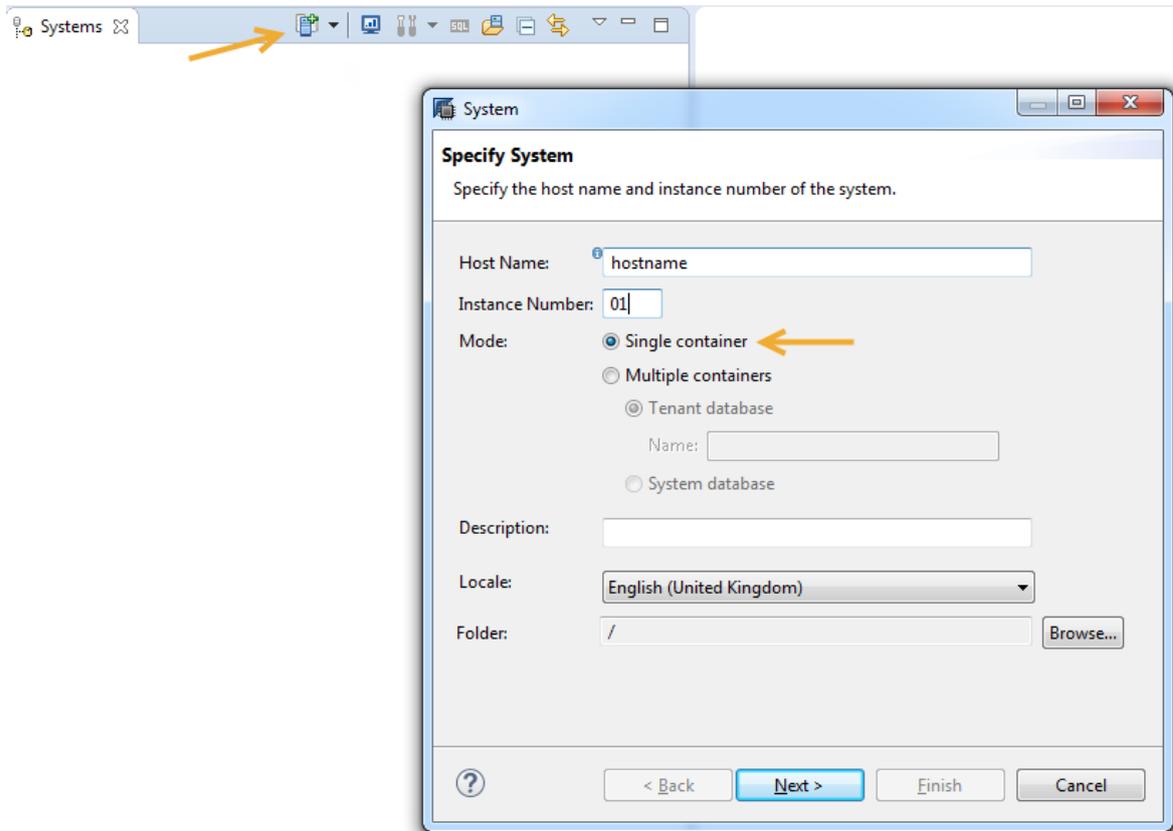
do you want to continue? (y/n): y
Checking installation...
Preparing package 'Saphostagent Setup'...
Preparing package 'Python Support'...
Preparing package 'Python Runtime'...
Preparing package 'Product Manifest'...
Preparing package 'Binaries'...
Preparing package 'Installer'...
Preparing package 'Ini Files'...
Preparing package 'hwCCT'...
Preparing package 'Emergency Support Package'...
Preparing package 'EPM'...
Preparing package 'Documentation'...
Preparing package 'Delivery Units'...
Preparing package 'DAT Languages'...
Preparing package 'DAT Configfiles'...
Creating System...
Extracting software...
Installing package 'Saphostagent Setup'...
Installing package 'Python Support'...
Installing package 'Python Runtime'...
Installing package 'Product Manifest'...
Installing package 'Binaries'...
Installing package 'Installer'...
Installing package 'Ini Files'...
Installing package 'hwCCT'...
Installing package 'Emergency Support Package'...
Installing package 'EPM'...
Installing package 'Documentation'...
Installing package 'Delivery Units'...
Installing package 'DAT Languages'...
Installing package 'DAT Configfiles'...
Creating instance...
  hdbparam: Working configuration directory: "/hana/shared/M00/global/hdb/custom/config"
  hdbnsutil: creating persistence ...
  hdbnsutil: writing initial topology...
  hdbnsutil: writing initial license: status check = 2
Starting SAP HANA Database system...
Starting 7 processes on host 'XXXXXXXXXX':
  Starting on 'XXXXXXXXXX': hdbcompilerver, hdbdaemon, hdbindexserver, hdbnameserver, hdbpreprocessor, hdbwebdispatcher, hdbxsengine
  Starting on 'XXXXXXXXXX': hdbcompilerver, hdbdaemon, hdbindexserver, hdbpreprocessor, hdbwebdispatcher, hdbxsengine
  Starting on 'XXXXXXXXXX': hdbcompilerver, hdbdaemon, hdbindexserver, hdbwebdispatcher, hdbxsengine
  Starting on 'XXXXXXXXXX': hdbdaemon, hdbindexserver, hdbwebdispatcher, hdbxsengine
  Starting on 'XXXXXXXXXX': hdbdaemon, hdbwebdispatcher, hdbxsengine
  All server processes started on host 'XXXXXXXXXX'.
Importing delivery units...
Importing delivery unit HCO_INA_SERVICE
Importing delivery unit HANA_DT_BASE
Importing delivery unit HANA_IDE_CORE
Importing delivery unit HANA_TALCONFIG
Importing delivery unit HANA_UI_INTEGRATION_SVC
Importing delivery unit HANA_UI_INTEGRATION_CONTENT
Importing delivery unit HANA_XS_BASE
Importing delivery unit HANA_XS_DBUTILS
Importing delivery unit HANA_XS_EDITOR
Importing delivery unit HANA_XS_IDE
Importing delivery unit HANA_XS_LM
Importing delivery unit HOC_ADMIN
Importing delivery unit HOC_IDE_CORE
Importing delivery unit HOC_SEC_CP
Importing delivery unit HOC_SEC_BASE
Importing delivery unit HOC_XS_LR
Importing delivery unit SARUIS_1
Importing delivery unit SAR_WATT
Importing delivery unit HANA_BACKUP
Importing delivery unit HANA_HDBLCH
Importing delivery unit HANA_SEC_BASE
Importing delivery unit HANA_SEC_CP
Importing delivery unit HANA_ADMIN
Installation done
Log file written to '/var/tmp/hdb_M00_install_2015-07-07_11.12.02/hdbinst.log' on host 'XXXXXXXXXX'.

```

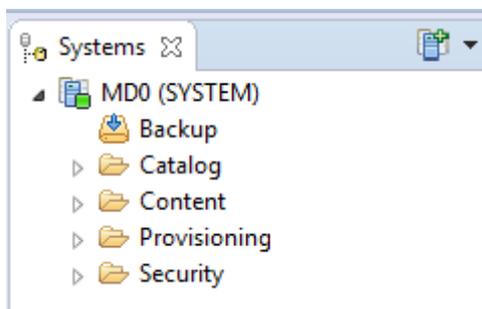
7. Check for errors in the installation log:

```
cat /var/tmp/hdb_<SID>_install_<timestamp>/hdbinst.log |grep ERROR
```

8. Open the SAP HANA studio and add the system:



The system MD0 appears in the *Systems* view:



## Results

The system MD0 is now up and running and you can install SAP DCC.

**Task overview:** [Tutorial: Migrating SAP DB Control Center to a Tenant Database \[page 196\]](#)

**Next task:** [Install and Configure SAP DCC on Single-Container System \[page 201\]](#)

---

## Related Information

[SAP HANA Server Installation and Update Guide \(HTML\)](#)

### 4.2.6.2 Install and Configure SAP DCC on Single-Container System

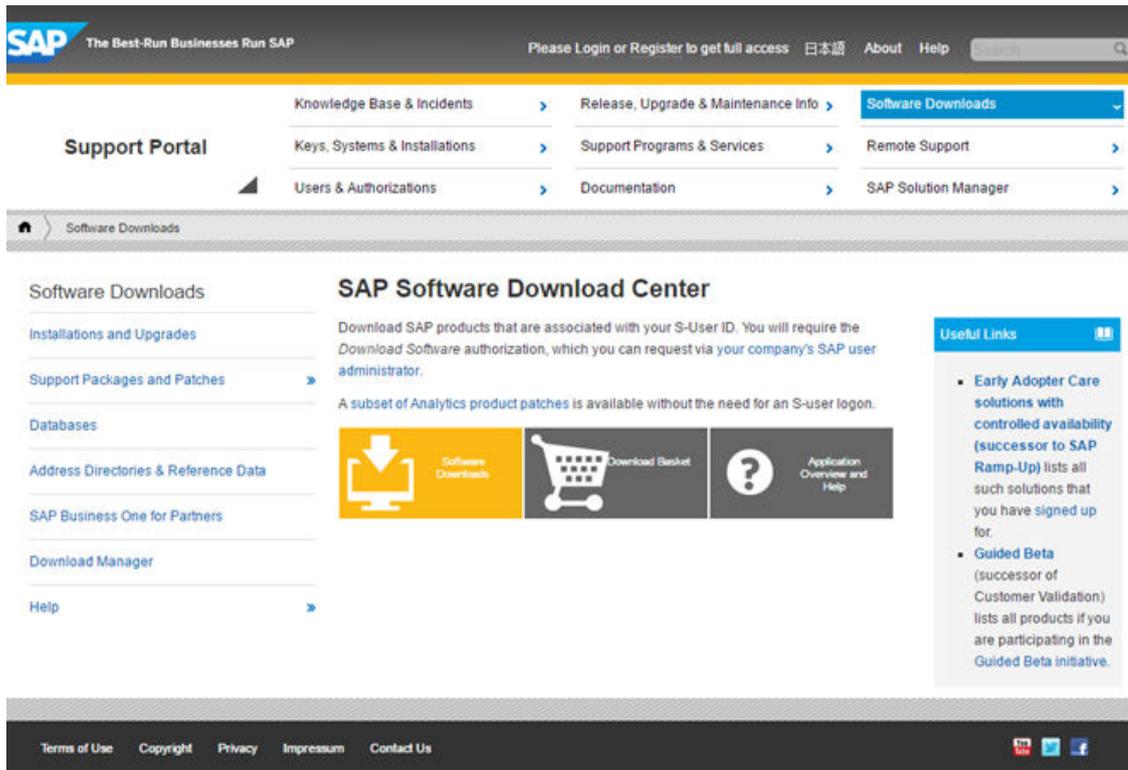
In this step, you'll install SAP DCC by importing the corresponding delivery unit (DU) into the newly installed system. Then, you'll configure SAP DCC and add systems to be monitored.

#### Prerequisites

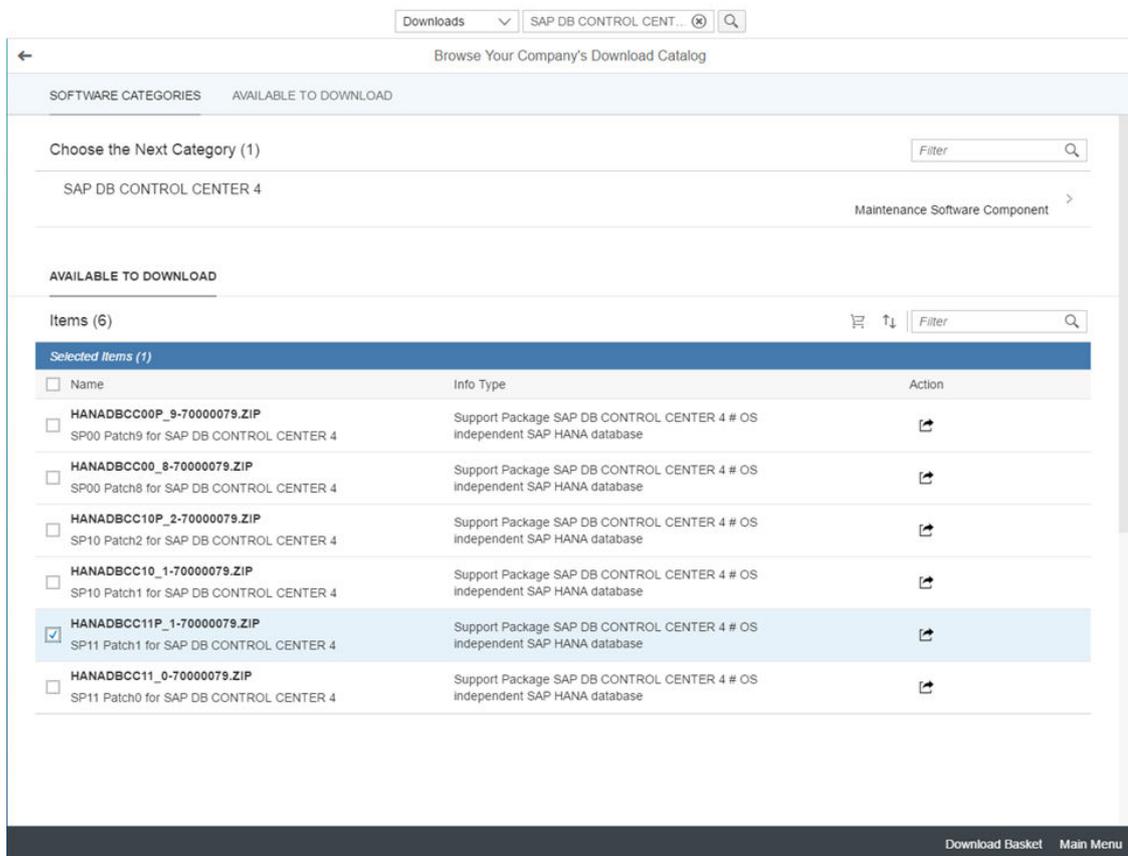
- You're familiar with deploying and configuring SAP DCC as described in the *SAP DB Control Center 4 Guide* and SAP Note 2188247.
- You're a registered user of SAP Service Marketplace and you have the Download Software authorization.
- Your database user has the following privileges required to import the SAP DCC DU using the SAP HANA studio:
  - System privilege REPO.IMPORT
  - Package privileges REPO.READ and REPO.ACTIVATE\_IMPORTED\_OBJECTS on the root package .REPO\_PACKAGE\_ROOT

#### Procedure

1. Download SAP DB Control Center 4 from the SAP Software Download Center (<http://support.sap.com/swdc>).
  - a. In the Software Download Center, click *Software Downloads*:



- b. From the SAP ONE Support Launchpad, open the *Software Downloads* app and search for **SAP DB Control Center**.

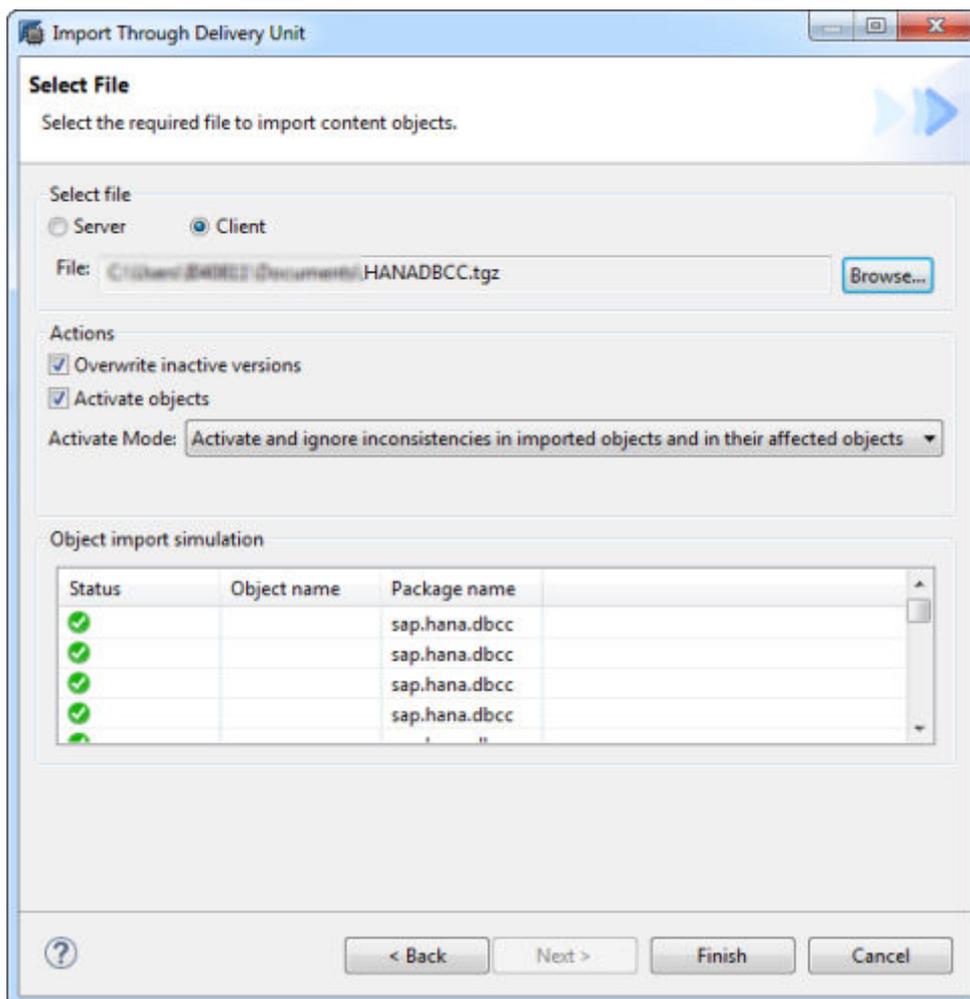


- c. Add the target release **SP11 Patch1 for SAP DB CONTROL CENTER 4** to the download basket, select the download basket, and then click the patch description to download.
2. Unpack the file into a temporary directory and note the location of the `HANADBCC.tgz` file.
3. Import the SAP DCC DU (`HANADBCC.tgz`) using the SAP HANA studio and check for successful activation.

### **i** Note

Instead of the SAP HANA studio, you could also use SAP HANA Application Lifecycle Management to import the DU. For more information, see *Import a Delivery Unit* in the *SAP HANA Application Lifecycle Management Guide*.

- a. In the SAP HANA studio, choose **File > Import > SAP HANA Content > Delivery Unit** and then *Next*.
- b. Select the target system you wish to import the SAP DCC DU to (in this example, MD0) and choose *Next*.
- c. Select the *Client* option and browse for the `HANADBCC.tgz` file that you unpacked above.
- d. Check the object import simulation to make sure there are no errors and choose *Finish* to complete the import.



- e. Review the information in the *Job Log* view to make sure the DU import completes successfully.

4. Set up technical user accounts in SAP DCC.

- a. In the SAP HANA studio, open the SQL console connected to the system MD0.
- b. Execute the SQL statements for configuring SAP DCC.

The SQL statements for configuring SAP DCC are available in the *Configuring with SQL* section of the *SAP DB Control Center 4 Guide*.

5. Log on to SAP DCC as the DCC configuration user (DCC\_CONFIG, by default).

The URL is: `http://<host>.corp:80<instance>/sap/hana/dbcc`.

6. On the launchpad, click *SAP DCC Setup*.

- a. If your environment requires a proxy server: in the Global Proxy section, click the *Enabled* box to configure SAP DCC to use a proxy server to connect to the systems it manages. Enter the fully qualified host name and port for the proxy server.

Configuring a default proxy activates the proxy options you can use when you add systems.

- b. Click *Next* to display the next configuration page.
- c. In the Administrator section, to designate the administrator who can add, import, and remove systems, do one of:
  - Click *Existing* if you want to administer SAP DCC with an account that already exists. Select the account from the drop-down list in the *Login* field.
  - Click *New* (or leave it selected) if you want to create a new account to administer SAP DCC. The account name defaults to DCC\_ADM. Enter and confirm a temporary password. The password must have a least eight characters and include at least one uppercase letter, at least one lowercase letter, and at least one digit.
- d. In the Collector section, enter and confirm a password for the technical user account SAP DCC will use to collect data. The password must have a least eight characters and include at least one uppercase letter, at least one lowercase letter, and at least one digit.

The account name defaults to DCC\_COLLECTOR.

The collector account runs collection jobs to gather the availability, performance, capacity, and alerts data that SAP DCC displays for each system it monitors. It also runs the jobs for multiple worker threads and cleaning up the message queue.

- e. (Optional) Select *Expiration Exempt* to prevent the collector password from expiring.

**i** Note

An expired password can disable the collector account's collection and housekeeping jobs. If you have password expiration policies in force, SAP recommends that you select *Expiration Exempt* for the collector account so that its password does not expire. For information on managing password lifetime, see the *SAP HANA Security Guide*.

- f. In the Technical section, enter and confirm a password for the technical user account SAP DCC uses for adding systems (including the SAP HANA host system) and for health monitoring. The password must have a least eight characters and include at least one uppercase letter, at least one lowercase letter, and at least one digit.

The account name defaults to SAPDBCC.

- g. (Optional) Select *Expiration Exempt* to prevent the technical user password from expiring.

### **i** Note

An expired password can disable monitoring functions. If you have password expiration policies in force, SAP recommends that you select *Expiration Exempt* for the technical user account so that its password does not expire. For information on managing password lifetime, see the *SAP HANA Security Guide*.

h. Click *Finish* to save your configuration changes.

7. Register the systems to be monitored in the system directory.

For more information about adding systems, see *System Directory* in the *SAP DB Control Center 4 Guide*.

8. Check the status of registered systems in the Enterprise Health Monitor.

For more information, see *Enterprise Health Monitor* in the *SAP DB Control Center 4 Guide*.

**Task overview:** [Tutorial: Migrating SAP DB Control Center to a Tenant Database \[page 196\]](#)

**Previous task:** [Install an SAP HANA System in Single-Container Mode \[page 197\]](#)

**Next task:** [Convert and Configure System for Multitenancy \[page 205\]](#)

## Related Information

[SAP DB Control Center 4 Guide \(HTML\)](#)

[SAP Note 2188366](#)

[Request User for SAP Service Marketplace](#)

### 4.2.6.3 Convert and Configure System for Multitenancy

In this step, you'll convert the MD0 system on which SAP DCC is running into a multiple-container system using the SAP HANA database lifecycle manager (HDBLCM). Afterward, SAP DCC will be running on the first tenant database in the multitenant system.

## Prerequisites

You have operating system access to the SAP HANA system.

## Procedure

### **i** Note

The following procedure describes how to convert your system using the Web user interface of the SAP HANA database lifecycle manager. 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.

1. Open the SAP HANA HDBLCM Web user interface by entering the following URL in an HTML5-enabled browser:

```
https://hostname:1129/lms1/HDBLCM/<SID>/index.html
```

### **i** Note

The URL is case sensitive. Make sure you enter upper and lower case letters correctly.

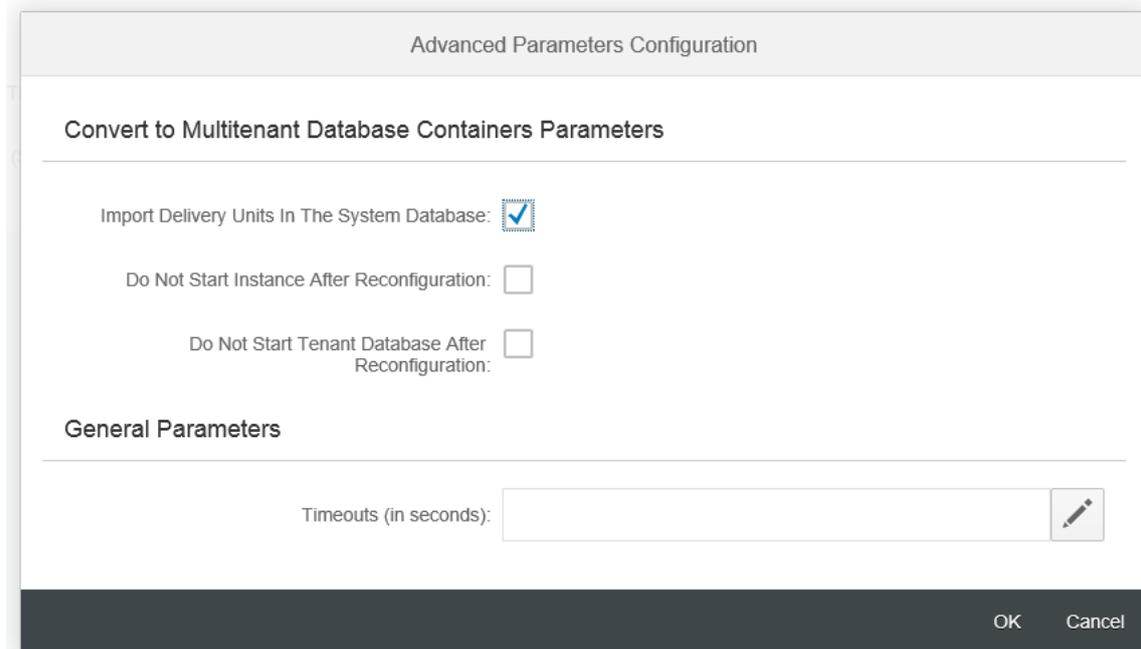
2. Click the tile *Convert to Multitenant Database Containers*.



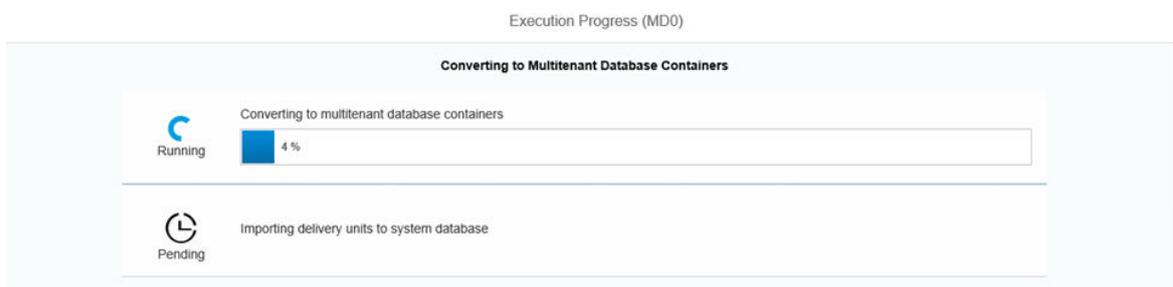
3. Enter the password of the `<sid>adm` user, specify the (new) password of the SYTSEM user of the system database (SYSTEMDB), and then click *Next*.

4. Open the advanced configuration dialog by clicking the personalization icon in the bottom left of the footer bar.

For the purposes of this tutorial, do not change the default configuration. In particular, it is important that all auto-content delivery units are imported into the system database.



5. Close the advanced configuration to return to the summary page. Then click *Run* to finalize the configuration and start the conversion.



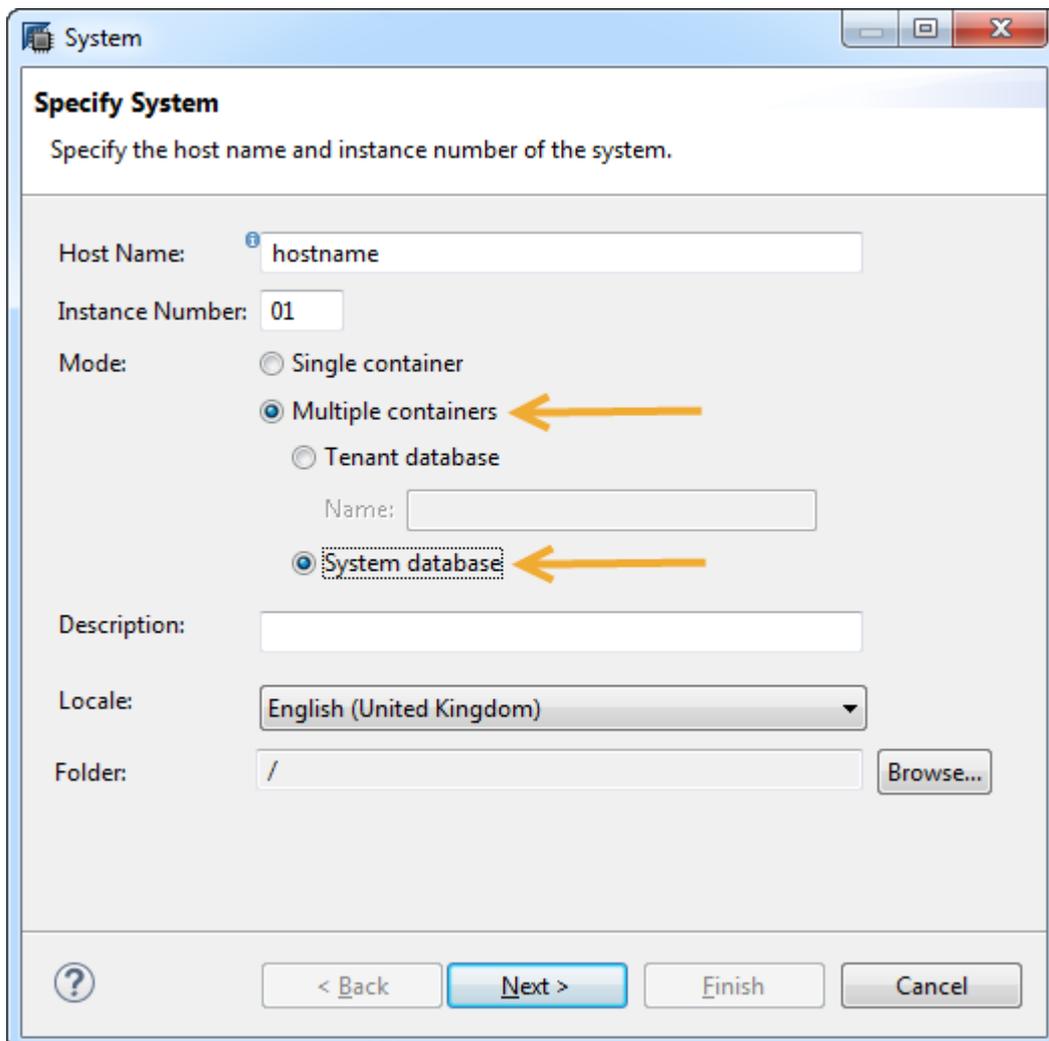
The system now supports multitenancy and has two databases: the system database (`SYSTEMDB`) and one tenant database (`MD0`) that corresponds to the original system.

### **i** Note

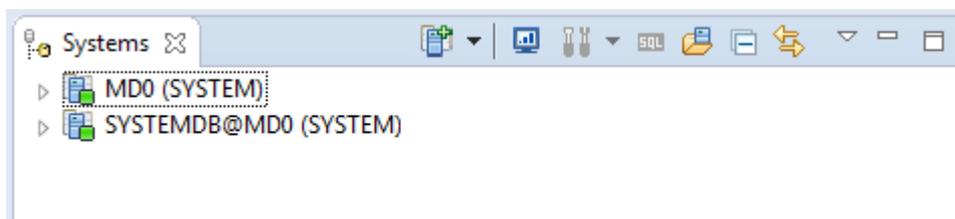
When you convert a system, the name of the first tenant database is the same as the SID of the system.



6. In the SAP HANA studio, add the system database with the SYSTEM user.

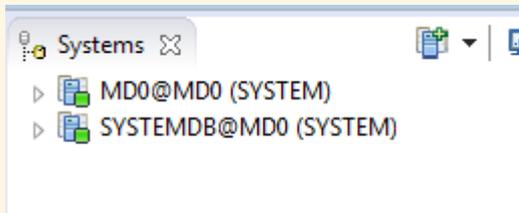


Now, if you refresh the *Systems* view, you'll see both the system database and the **tenant database** MD0 as up and running.



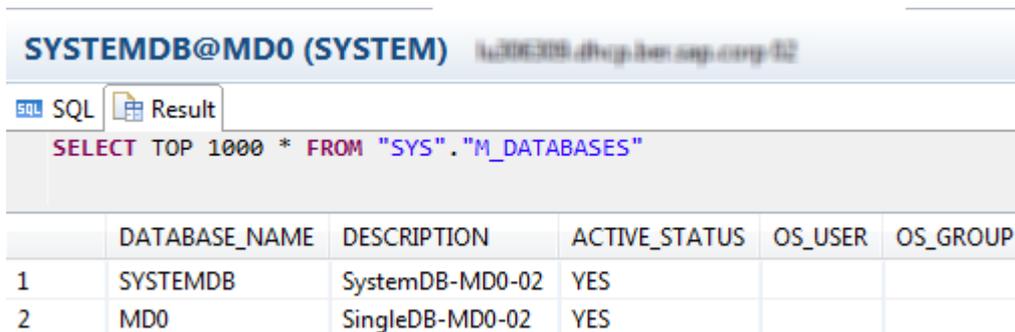
### Note

Because you added MD0 when it was a single-container system, it still looks like a single-container system. If you were to remove it and re-add it as a tenant database, then it would be identified as a tenant database (MD0@MD0):



7. Connected to the system database, check the `M_DATABASES` view to confirm the creation of the system database and tenant database:

```
SELECT TOP 1000 * FROM "SYS"."M_DATABASES";
```

A screenshot of the SQL console interface. The title bar shows 'SYSTEMDB@MD0 (SYSTEM)'. Below the title bar, there are tabs for 'SQL' and 'Result'. The 'Result' tab is active, showing the query 'SELECT TOP 1000 \* FROM "SYS"."M\_DATABASES"'. Below the query, there is a table with the following data:

	DATABASE_NAME	DESCRIPTION	ACTIVE_STATUS	OS_USER	OS_GROUP
1	SYSTEMDB	SystemDB-MD0-02	YES		
2	MD0	SingleDB-MD0-02	YES		

8. Configure HTTP access to the tenant database by entering the URLs by which the tenant database is publicly accessible in the `xsengine.ini` file.

You can do this in the Administration editor on the *Configuration* tab, or by executing the following statements in the SQL console (connected to the system database):

```
ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', 'MD0) SET ('public_urls', 'http_url') = 'http://<tenant_FDQN>:8000 ' WITH RECONFIGURE;  
ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', 'MD0) SET ('public_urls', 'https_url') = 'https://<tenant_FDQN>:4300 ' WITH RECONFIGURE;
```

9. Check the status of the `xsengine` service on the tenant database by opening the Administration editor of the tenant database and choosing the *Landscape* tab.

If the `xsengine` service is in the list of services, the XS server on the tenant database is running as a separate service. After a conversion, this should be the case. However, if you want it to run as an embedded service in the index server, follow these steps:

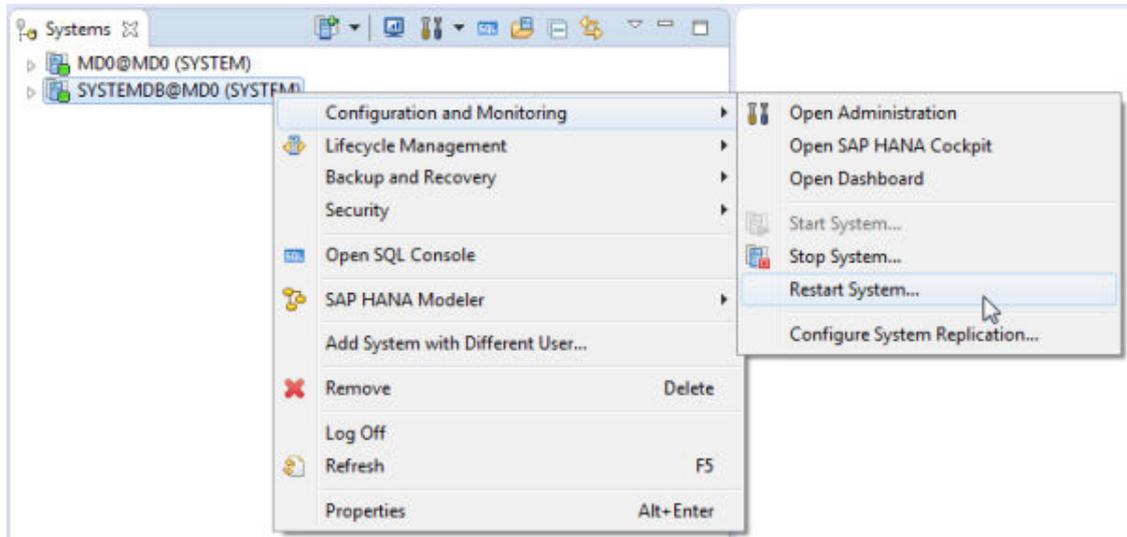
- a. Open the Administration editor of the system database and choose the *Configuration* tab.
- b. In the `xsengine.ini` change the value of the `[httpserver] embedded` parameter to **true**.
- c. Remove the existing `xsengine` server by executing the following statement in the SQL console (connected to the system database):

```
ALTER DATABASE MD0 REMOVE 'xsengine' AT '<host>:<port>';
```

➔ Tip

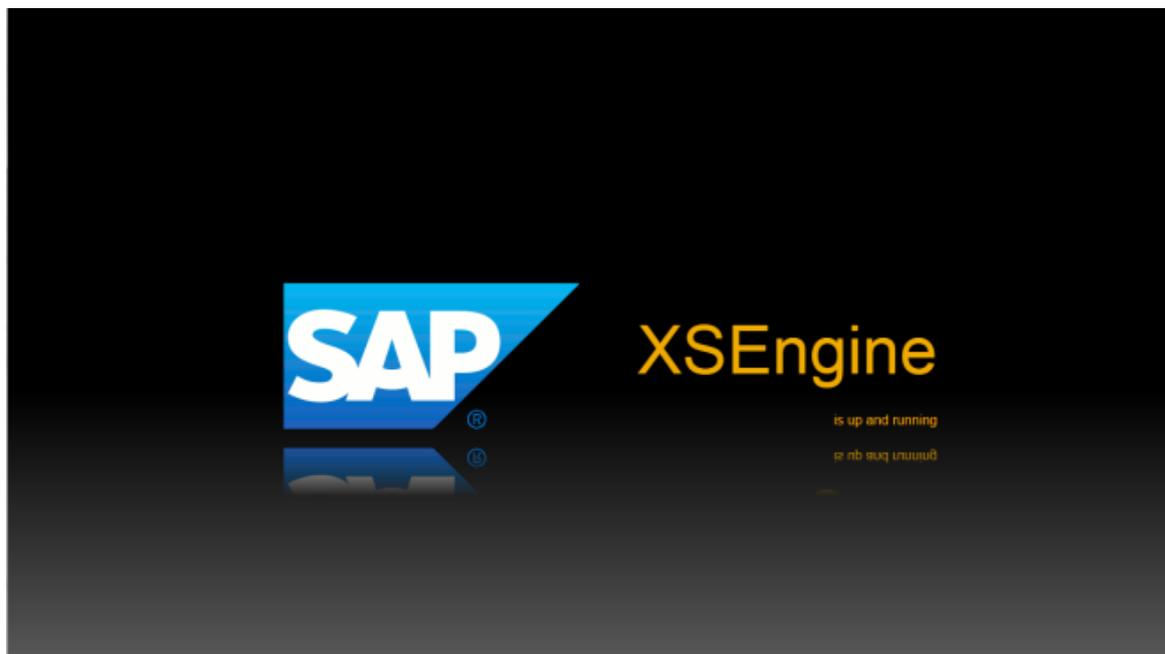
You can see the port of the xsengine service on the *Landscape* tab.

- d. Restart the system.



10. Configure the virtual host names by which the tenant database is publicly accessible in your Domain Name Sever (DNS).
11. Open a Web browser and check that the XS servers of both the system database and the tenant database are accessible.
  - The system database is accessible under the localhost: `https://<localhost>:80<instance>`, or if SSL is configured: `https://<localhost>:4300`.
  - The tenant database MD0 is accessible under the full qualified domain name (FQDN) that you specified in step 7: `http://<tenant_FQDN>:80<instance>` or if SSL is configured `https://<tenant_FQDN>:43<instance>`.

In both cases you should see the following screen:



12. Check that SAP DCC has been successfully migrated to the tenant database and is accessible by logging on to SAP DCC on the tenant database:

`http://<tenant_FQDN>:80<instance>/sap/hana/dbcc`

13. Change the connection information of the system monitored by SAP DCC.

If you previously configured SAP DCC to monitor the SAP HANA system on which it is installed, the connectivity information will now be incorrect as SAP DCC has moved to the tenant database. The only workaround is to remove the old system reference and to re-register it to SAP DCC with the new connection information.

- a. Log on to SAP DCC as a user with the `DBCCAdmin` role.
  - b. Open the system directory.
  - c. Select the system representing the single-container SAP HANA system on which SAP DCC was originally installed.
  - d. On the *System Information* screen, select *Remove System* and confirm.
  - e. Add the SAP HANA system representing the tenant database.
14. Optional: Register the system database in SAP DCC.

➔ Remember

The host of the system database is not the FQDN but the localhost name.

**Task overview:** [Tutorial: Migrating SAP DB Control Center to a Tenant Database \[page 196\]](#)

**Previous task:** [Install and Configure SAP DCC on Single-Container System \[page 201\]](#)

## Related Information

[Converting an SAP HANA System to Support Multitenant Database Containers \[page 572\]](#)

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

### 4.3 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, hosts, and services.

SAP HANA's configuration files contain parameters for global system configuration (`global.ini`), as well as the configuration of each service in the system (for example, `indexserver.ini`). You can access and edit configuration files on the *Configuration* tab of the Administration editor. They are stored at the following default location on the server: `/hana/shared/$SID/global/hdb/custom/config`.

Configuration files are stored on the SAP HANA server at the following location: `<sapmnt>/<SID>/global/hdb/custom/config/DB_<dbname>`. By default, `<sapmnt>` is `/hana/shared`.

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 If a system-specific value is not configured for a property, the default value applies.
Host	The host-specific value for the property 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.

#### **i** Note

A further layer ("database") is available in systems with multitenant database containers. For more information, see *System Properties in Multitenant Database Containers*.

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.

#### **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.

### ➔ Tip

To alter parameters of the secondary site in a system replication scenario, note that you cannot do this from SAP HANA studio. Instead you alter the parameters directly in the .ini files on the secondary site. Afterwards reconfigure the database using `hdbnsutil -reconfig`.

## Related Information

[Configuration Parameters in Multiple-Container Systems \[page 213\]](#)

[SAP Note 203611](#)

### 4.3.1 Configuration Parameters in Multiple-Container Systems

In addition to the layers "default", "system", and "host", system configuration files in multiple-container systems have a "database" layer to facilitate the configuration of properties for individual databases.

#### Database-Specific Configuration

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`
- `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. In addition, it allows client applications designed for single-container systems (where the system and the database **are** the same thing) to query the system-layer value and still retrieve the correct value for the tenant database. 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.

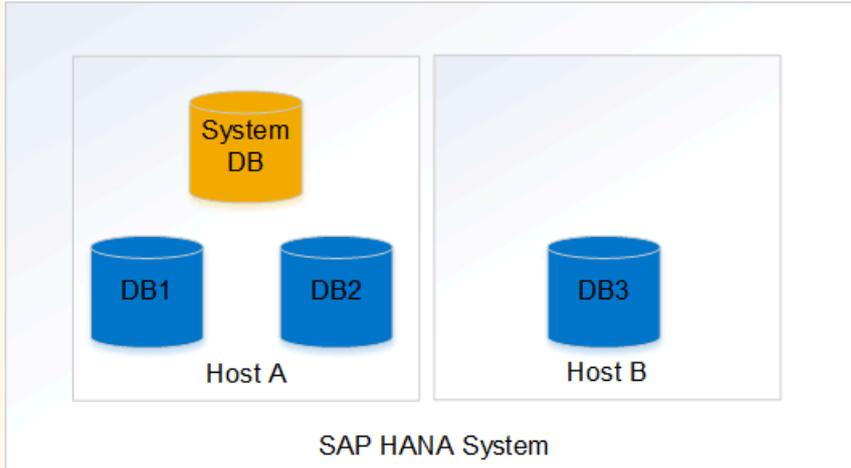


You can view actual configuration values in the Administration editor of the SAP HANA studio ([Configuration](#) tab) or by querying the following system views:

- 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 multiple-container 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 ← Effective value
indexserver.ini	DATABASE			execution	max_concurrency	20

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 ← Effective value
indexserver.ini	SYSTEM			execution	max_concurrency	20
indexserver.ini	DATABASE			execution	max_concurrency	20

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 ← Effective value
indexserver.ini	SYSTEM			execution	max_concurrency	20

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 203611](#)

## 4.3.2 Change a System Property

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
<b>System</b>	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.
<b>Host</b>	<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>
<b>Database</b>	<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>

### **i** 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.

### 4.3.3 Reset a System Property

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.

- To delete a user-defined value and restore the default value, you can choose one of the following methods:

Procedure	Result
Delete with automatic reset: <ol style="list-style-type: none"> <li>In the context menu of the configuration parameter, choose <i>Delete</i>. The <i>Delete Configuration Value</i> dialog box appears.</li> <li>Choose the layer whose user-defined values you want to delete.</li> <li>Choose <i>Delete</i>.</li> </ol>	The user-defined value(s) are cleared and the default value(s) are re-applied. <div style="background-color: #fff9c4; padding: 5px; margin-top: 10px;"> <p><b>i Note</b> If you added a new parameter to a section, <i>Delete</i> deletes the parameter.</p> </div>
Manually restore default: <ol style="list-style-type: none"> <li>In the context menu of the configuration parameter, choose <i>Change...</i> The <i>Change Configuration Value</i> dialog box appears.</li> <li>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>.</li> <li>Choose <i>Save</i>.</li> </ol>	The user-defined value(s) are cleared and the default value(s) are re-applied.

## 4.3.4 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

- In the Administration editor, choose the *Configuration* tab.
- Navigate to the `indexserver.ini` file and expand the `session` section.
- Configure the `reserved_connections` parameter by specifying the number of connections you want to reserve.  
The default number of reserved connections is 10. The minimum number is 1.
- 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.

### 4.3.5 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 studio can use this property to alter behavior.

#### Prerequisites

You have the system privilege INFILE ADMIN.

#### Procedure

1. In the Administration editor, choose the *Configuration* tab
2. Navigate to the `global.ini` file and expand the `system_information` section.
3. Configure the `usage` parameter.

Although you can enter any value, the values listed below can be used by clients to alter behavior. For example, the SAP HANA studio evaluates this parameter to warn 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). Note that there is no change in the behavior of the SAP HANA studio for other values.

- production
  - test
  - development
  - custom (default)
4. Restart the system.

## 4.4 Managing SAP HANA Licenses

License keys are required to use SAP HANA databases. You can install, and delete license keys using the SAP HANA studio, the SAP HANA HDBSQL command line tool, and SQL.

### Related Information

[License Keys \[page 222\]](#)

[Check the Current License Key \[page 224\]](#)

[Install a Permanent License \[page 224\]](#)

[Delete an Existing Permanent License Key \[page 226\]](#)

### 4.4.1 License Keys

License keys are required to use SAP HANA. SAP HANA supports two kinds of license key: temporary license keys and permanent license keys.

While temporary license keys are automatically installed in an SAP HANA system, permanent license keys have to be requested on the SAP Service Marketplace and applied to the individual SAP HANA system.

#### **i** Note

Systems that support multitenant database containers require a single license key, regardless of the number of tenant databases.

### Temporary License Keys

A temporary license key, which is valid for 90 days, is automatically installed with a new SAP HANA system. During this period, you should request and apply a permanent license key.

### Permanent License Keys

You can request a permanent license key on the SAP Service Marketplace under [Keys & Requests](#). Permanent license keys are valid until the predefined expiration date. Furthermore, they specify the amount of memory licensed to the target SAP HANA installation. Before a permanent license key expires, you should request and apply a new permanent license key. If a permanent license key expires, a temporary license key valid for 28 days is automatically installed. During this time, you can request and install a new permanent license key.

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 system 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

### **i** Note

Although enforced license keys currently only apply to SAP Business One, it is technically possible to install such a license in an SAP HANA instance 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, no memory consumption check is enforced. 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.

## System Lockdown

The system goes into lockdown mode in the following situations:

- The temporary license key has expired.
- You were using a temporary license key and the hardware key has changed.
- The permanent license key has expired and you did not renew it within 28 days.
- 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 your database.

In lockdown mode, no queries are possible. 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.

In addition, the database cannot be backed up in lockdown mode.

### **i** Note

If a system has locked down due to an invalid or expired license, the icon indicating the operational status of the system in the *Systems* view and the *System Monitor* changes accordingly.

## Related Information

[SAP Service Marketplace](#) 

## 4.4.2 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.

## 4.4.3 Install a Permanent License

To use SAP HANA, you must request and install a valid permanent license key.

### 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 a 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*.

Permanent licenses are sent as e-mail attachments.

3. Install the license key using one of the tools below.

### **i** Note

If the system supports multitenant database containers, you must install the license in the system database.

Tool	Steps
SAP HANA studio	<ol style="list-style-type: none"> <li>1. In the <i>Systems</i> view, right-click the system and choose ► <i>Properties</i> ► <i>License</i> ▾.</li> <li>2. In the <i>Request License Key</i> area of the <i>System License</i> page, choose <i>Install License Key</i> and select the file that you received by e-mail.</li> </ol>
SQL console	Execute the SQL statement <code>SET SYSTEM LICENSE '&lt;license file content&gt;'</code>
HDBSQL	<h3><b>i</b> Note</h3> <p>If using HDBSQL, you need to enable multiple-line mode before executing the statement above. For more information, see <i>Execute Long Commands in Multiple-Line Mode</i>.</p>

### **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.

## Related Information

[SAP Support Portal](#)

[SAP HANA HDBSQL \(Command-Line Reference\) \[page 1346\]](#)

[Run Long Commands in Multiple-Line Mode \[page 1356\]](#)

### 4.4.4 Delete an Existing Permanent License Key

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

Uninstall the license key using one of the following options:

Option	Steps
SAP HANA Studio	<ol style="list-style-type: none"><li>1. In the <i>Systems</i> view, right-click the system and choose <b>► Properties ► License ►</b>.</li><li>2. On the <i>System License</i> page choose <i>Delete License Key</i>.</li></ol>
SQL console	Execute the SQL command <code>UNSET SYSTEM LICENSE ALL</code>
HDBSQL	

#### 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 \[page 65\]](#)

[SAP HANA HDBSQL \(Command-Line Reference\) \[page 1346\]](#)

---

## 4.5 Monitoring the SAP HANA Database

It's important that you monitor the operation of the SAP HANA database 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.

### Related Information

[Monitoring in SAP HANA Studio \[page 227\]](#)

[Monitoring in SAP HANA Cockpit \[page 290\]](#)

### 4.5.1 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.

In addition, you can access the SAP HANA cockpit

### Related Information

[System Monitor \[page 228\]](#)

[Administration Editor \[page 229\]](#)

[Open SAP HANA Cockpit from SAP HANA Studio \[page 25\]](#)

## 4.5.1.1 System Monitor

The System Monitor 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.

### **i** 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...	52,574/4031,73	49,66/4031,73	0,21/4031,73	132,75/960,13	157,80/960,13	36,03/895,35	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	18,25/87,15	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	70,50/485,13	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
UTO (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

[Create and Authorize a User \[page 704\]](#)

[Options for Customizing the System Monitor \[page 229\]](#)

## 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

### 4.5.1.1.1 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 System Monitor, if for example you have a very large number of systems
 (Properties)	Allows you to configure properties of the System Monitor, 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

### 4.5.1.2 Administration Editor

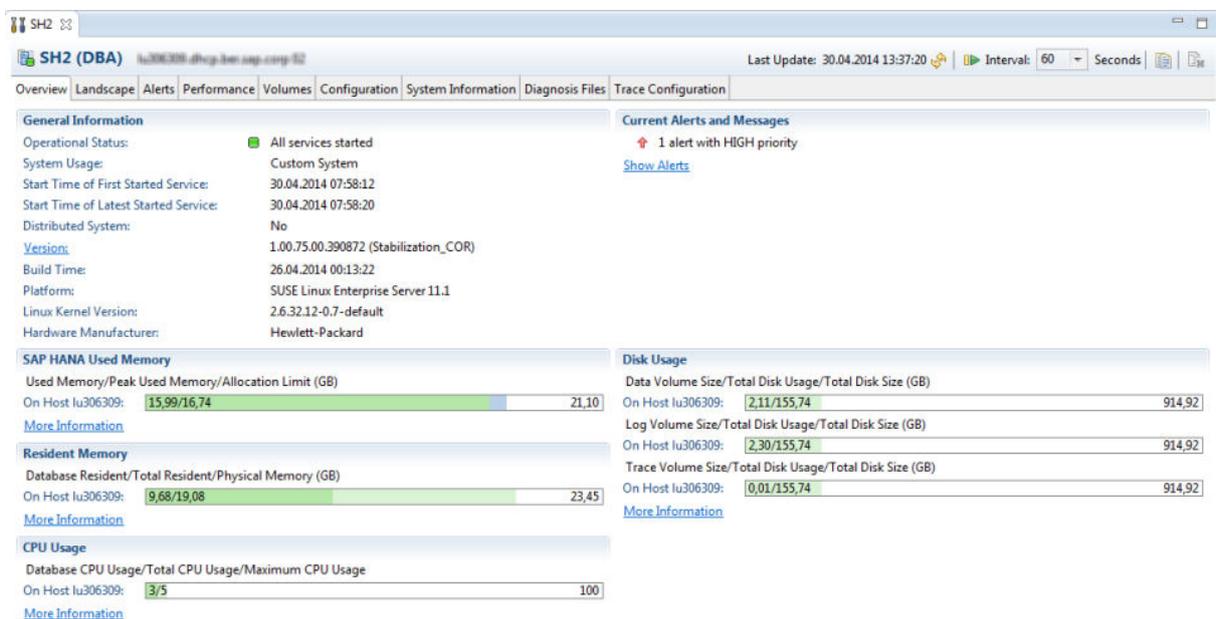
To identify problems early and avoid disruptions, you need to monitor your systems continuously. While the System Monitor provides you with an overview of all systems at a glance, the Administration editor allows you to drill down into the details of resource usage and performance for each system.

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

## **i** 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 \[page 704\]](#)

### 4.5.1.3 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.

## 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 (*SAP HANA Studio Displays System Status as Yellow*).

## Related Information

[Monitoring SAP HANA Systems During Stop and Start \[page 103\]](#)

[SAP Note 1639568](#)

### 4.5.1.3.1 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.

### 4.5.1.4 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"> <li>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</li> <li>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 <a href="#">Landscape &gt; Secondary System Status</a> tab.</li> </ul>
Alerts and Messages	Priority-rated alerts and messages reported by the system
Database Used Memory	<p>The following key indicators of memory usage are displayed:</p> <ul style="list-style-type: none"> <li><a href="#">Used Memory</a> The total amount of memory currently in use by SAP HANA is referred to as its used memory.</li> <li><a href="#">Peak Used Memory</a> The <a href="#">Used Memory</a> value is a current measurement. The <a href="#">Peak Used Memory</a> value is the highest used memory value recorded. This is useful for keeping track of the maximum value for used memory over time.</li> </ul> <div style="background-color: #fff9c4; padding: 10px; margin: 10px 0;"> <p><b>i Note</b></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 <a href="#">Landscape &gt; Services</a> tab. From the context menu, choose <a href="#">Reset Memory Statistics</a>.</p> </div> <ul style="list-style-type: none"> <li><a href="#">Allocation Limit</a> 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.</li> </ul>
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 <a href="#">Used Memory</a> value is lower than the <a href="#">Database Resident</a> 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.

Screen Area	Information Available
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

[Monitoring Memory Usage \[page 270\]](#)

### 4.5.1.5 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	<p>Indicates the status of the service</p> <p>The following statuses are possible:</p> <ul style="list-style-type: none"> <li>• The service is started (  )</li> <li>• The service is stopping or starting (  )</li> <li>• The service is stopped (  )</li> </ul> <p>The daemon service displays the  icon while the host or any of its services are starting or stopping.</p>
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
Detail	<p>Role of the host on which the service is running</p> <ul style="list-style-type: none"> <li>• Master The host is the active master host.</li> <li>• &lt;Empty&gt; The host is an active slave host.</li> <li>• Standby The host is a standby host.</li> </ul> <p>This is relevant only for distributed systems. For more detailed information, see the <a href="#">Hosts</a> sub-tab.</p>
Start Time	<p>Time at which the service started</p> <p>Two of the times in this column should match the <a href="#">Start time of first started service</a> and <a href="#">Start time of latest started service</a> shown on <a href="#">Overview</a> 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

[Monitoring Memory Usage \[page 270\]](#)

## 4.5.1.6 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.

### 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"> <li>• YES (  ) All services are active.</li> <li>• PARTIAL (  ) Some services active.</li> <li>• STARTING (  ) Some services are active, some are starting.</li> <li>• STOPPING (  ) Some services are active, some are stopping</li> <li>• NO (  ) No services active</li> </ul>
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"> <li>• OK (  ) The system is operational and the host's actual role corresponds to its configured role.</li> <li>• IGNORE (  ) The system is operational. The host is configured as a standby host and is available, but not in use.</li> <li>• INFO (  ) The system is operational. The host's actual role is different from its configured role.</li> <li>• WARNING (  ) The system is not operational. The host will become available after start-up or failover.</li> <li>• ERROR (  ) The system is not operational. The host is missing.</li> </ul>

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"> <li>• Empty Failover is neither active nor pending.</li> <li>• WAITING ... SEC The host has failed. The system is waiting to fail over.</li> <li>• WAITING The host has failed. The system is waiting for the host to restart to prevent unnecessary fail-over.</li> <li>• FAILOVER TO &lt;host&gt; The host has failed and failover to a target host is in progress.</li> <li>• FAILBACK TO &lt;host&gt; 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.</li> <li>• FAILED Failover is not possible, for example, no further standby hosts available. For more information, see the nameserver trace.</li> </ul>
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"> <li>• 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.</li> <li>• SLAVE Additional hosts in your system are configured as slave name servers. The configured nameserver role of these hosts is SLAVE.</li> </ul>
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"> <li>• 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.</li> <li>• SLAVE The actual nameserver role of the remaining hosts configured as master and slave hosts is SLAVE.</li> </ul>

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"> <li>• WORKER</li> <li>• STANDBY</li> </ul> <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"> <li>• 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).</li> <li>• 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.</li> <li>• 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.</li> </ul> <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="background-color: #fff9c4; padding: 10px;"> <p><b>i Note</b></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>

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"> <li>• WORKER Worker host for database processing</li> <li>• STANDBY Standby host for database processing</li> </ul> <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"> <li>• EXTENDED_STORAGE_WORKER Worker host for SAP HANA dynamic tiering</li> <li>• EXTENDED_STORAGE_STANDBY Standby host for SAP HANA dynamic tiering</li> <li>• ETS_WORKER Worker host for SAP HANA accelerator for SAP ASE</li> <li>• ETS_STANDBY Standby host for SAP HANA accelerator for SAP ASE</li> <li>• RDSYNC Host for SAP HANA remote data sync</li> <li>• STREAMING Host for SAP HANA smart data streaming</li> </ul> <div style="background-color: #fff9c4; padding: 10px; margin: 10px 0;"> <p><b>⚠ Caution</b></p> <p>Be aware that you need additional licenses for SAP HANA options and capabilities. For more information, see <a href="#">Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities [page 1360]</a>.</p> </div> <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"> <li>• XS_STANDBY Standby host for SAP HANA XS advanced runtime</li> <li>• XS_WORKER Host for SAP HANA XS advanced runtime</li> </ul> <div style="background-color: #fff9c4; padding: 10px; margin: 10px 0;"> <p><b>i Note</b></p> <p>Multiple host roles are <b>not</b> supported in production environments. However, if XS advanced runtime is installed, hosts can share multiple roles.</p> </div> <div style="background-color: #fff9c4; padding: 10px; margin: 10px 0;"> <p><b>i Note</b></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,</p> </div>

Column	Description
	<p>please note that it is not recommended to immediately start migrating existing applications to the new capabilities.</p>
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"> <li>• <code>/usr/sap/&lt;SID&gt;/SYS/global/hdb/data</code> for data</li> <li>• <code>/usr/sap/&lt;SID&gt;/SYS/global/hdb/log</code> for logs</li> </ul> <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 your 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>The following statuses are possible:</p> <ul style="list-style-type: none"> <li>• <code>&lt;Empty&gt;</code> Host has not been marked for removal.</li> <li>• REORG PENDING A redistribution operation is required to move tables to other hosts.</li> <li>• 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>.</li> <li>• REORG FAILED A redistribution operation was executed and failed. For more information, query the system table <code>SYS.REORG_STEPS</code>.</li> <li>• REORG FINISHED A redistribution operation has completed. The host can be uninstalled.</li> <li>• REORG NOT REQUIRED A redistribution operation not required. The host can be uninstalled.</li> </ul>

## Related Information

[High Availability for SAP HANA \[page 774\]](#)

[Redistribute Tables Before Removing a Host \[page 1009\]](#)

## 4.5.1.7 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

### 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 \[page 243\]](#)

[The Statistics Service \[page 245\]](#)

[Check Information \[page 244\]](#)

### 4.5.1.7.1 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

## 4.5.1.7.2 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 <a href="#">Next Run</a> 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 <a href="#">On Schedule</a> column. If a check is not on schedule ( <i>No</i> in the <a href="#">On Schedule</a> 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.

### **i** Note

These tables are only available if you have migrated to the new implementation of the statistics server available with SPS 07. For more information, see SAP Note 1917938.

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.  <b>i</b> Note It also contains the scheduling information for historical data collection.

## Related Information

[SAP Note 1917938](#)

### 4.5.1.7.3 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 \[page 245\]](#)
- [Technical Implementation \[page 246\]](#)
- [Statistics Server or Statistics Service? \[page 246\]](#)
- [Data Management in the Statistics Service \[page 247\]](#)
- [The Statistics Service in Multitenant Database Containers \[page 248\]](#)

## 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, depending on your system set-up, the statistics service or the statistics server.

The statistics service (or statistics server) 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 the SAP HANA studio, SAP HANA cockpit, or SAP DB Control Center.

### **i** Note

The SAP HANA cockpit requires the statistics service to access alert and check information.

## Technical Implementation

The monitoring and alerting features of the SAP HANA database can be performed by two mechanisms: the statistics server or the statistics service.

The **statistics server**, which runs as a separate server process, is an enhanced index server with a monitoring extension on top. Data is collected remotely from the other servers and stored in measurement tables in the statistics server.

The **statistics service** is a simplified, more flexible implementation of the statistics server. It 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).

## Statistics Server or Statistics Service?

The embedded statistics service is the recommended mechanism because it offers more flexible configuration options and uses fewer system resources.

As of revision 93, the embedded statistics service is enabled by default after installation or upgrade. For more information, see SAP Note 2091313. In earlier revisions (but at least revision 74), the statistics server is the

default mechanism. For more information about how to migrate from the statistics server to the statistics service, see SAP Note 1917938.

If the statistics server component is deactivated in your system and the statistics service within the index server is active, the following system properties will be configured:

- `statisticserver.ini`: The property `[statisticsserver_general] active` has the value `false`
- `nameserver.ini`: The property `[statisticsserver] active` has the value `true`

## 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. It is not possible to add the statistics server as separate server. 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

[SAP Note 1917938](#)

[SAP Note 2091313](#)

[Multitenant Database Containers \[page 15\]](#)

### 4.5.1.7.4 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 server 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.

How the system responds to a failing check depends on whether the statistic server or statistics service is running.

Alerting Mechanism	Response
Statistics server	If a check fails to execute 10 times in succession, it is disabled for 6 hours. Then it is automatically re-enabled.

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"> <li>• INTERVALLENGTH</li> <li>• SKIP_INTERVAL_ON_DISABLE</li> </ul> <p>Once <math>INTERVALLENGTH \times SKIP\_INTERVAL\_ON\_DISABLE</math> 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. For more information about how to do this, see SAP Note 1991615.</p>

### **i** Note

The behavior described above also applies to the data collectors of the statistics server or service.

## Related Information

[SAP Note 1991615](#)

### 4.5.1.7.5 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

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. 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 server/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.

## 4.5.1.7.6 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.

#### 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.
- Choose the *Configure Check Thresholds* tab.
- 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.

- 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.

## 4.5.1.7.7 Configure Start Times of Periodic Checks

Some statistics server checks are performed every 6 or 24 hours. For example, the age of the last data backup is checked every 24 hours. You can configure the start times for these checks.

### Prerequisites

- You have the system privilege CATALOG READ and the SELECT privilege on the `_SYS_STATISTICS` schema.

#### Note

Both of these privileges are included in the role MONITORING.

- You have the system privilege INIFILE ADMIN.

### Procedure

#### Note

It is only possible to configure start times of periodic checks if the statistics server component is still available in your system, that is, you have not migrated to the new mechanism available with SPS 07.

- 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.
- Choose the *Configure Start Time on Check Intervals* tab.  
This tab has two sections, one for setting the start time of checks performed every 6 hours and one for setting the start time of checks performed every 24 hours. You can also see which checks are performed at these intervals.

#### Note

The intervals between checks cannot be changed.

- Set the start time for both periodic check types.
- Choose *OK*.

### Results

The start time for the checks is changed.

## 4.5.1.8 Monitoring System Performance

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.

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 \[page 253\]](#)

### 4.5.1.8.1 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  ( <a href="#">Configure Viewer</a> ) 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"><li>• <a href="#">Threads</a></li><li>• <a href="#">Sessions</a></li><li>• <a href="#">Blocked Transactions</a></li></ul>
Column Filter	Use this filter to see information by distinct values in visible columns. To open the column filter, choose  <a href="#">Filters...</a> in the toolbar on the top-right of the screen.	<ul style="list-style-type: none"><li>• <a href="#">Sessions</a></li><li>• <a href="#">SQL Plan Cache</a></li><li>• <a href="#">Expensive Statements</a></li><li>• <a href="#">Job Progress</a></li></ul>
Host/Service/Thread Type	Use these filters to show threads from a specific host or service, or of a specific type.	<a href="#">Threads</a>

Filter	Description	Sub-Tab
Free-text filter	Use this filter to find a specific character string or expression in the displayed information.	<ul style="list-style-type: none"> <li>• <a href="#">Sessions</a></li> <li>• <a href="#">Blocked Transactions</a></li> <li>• <a href="#">SQL Plan Cache</a></li> <li>• <a href="#">Expensive Statements</a></li> <li>• <a href="#">Job Progress</a></li> </ul>

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](#)

## 4.5.1.8.2 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.

### Note

With revision 56, `PlanExecutor` threads were replaced by `JobWorker` threads.

### 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.

### 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 \[page 469\]](#)

[Performance Trace \[page 474\]](#)

[Expensive Statements Trace \[page 475\]](#)

## 4.5.1.8.3 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).

#### 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 \[page 469\]](#)

[Performance Trace \[page 474\]](#)

[Expensive Statements Trace \[page 475\]](#)

### 4.5.1.8.4 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 \[page 469\]](#)

[Performance Trace \[page 474\]](#)

[Expensive Statements Trace \[page 475\]](#)

### 4.5.1.8.5 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.

## 4.5.1.8.6 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 \[page 475\]](#)

### 4.5.1.8.7 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.

### 4.5.1.8.8 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.

### 4.5.1.9 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.

Column	Description
Service/Volume	The service host and internal port You can expand the host/port to can 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

## 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"><li>• 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.</li><li>• 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>.</li><li>• Trace files Used size is zero for unused trace files and equals total size for used trace files.</li></ul>
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>

## 4.5.1.9.1 Persistent Data Storage in the SAP HANA Database

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_logvolmes` 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.

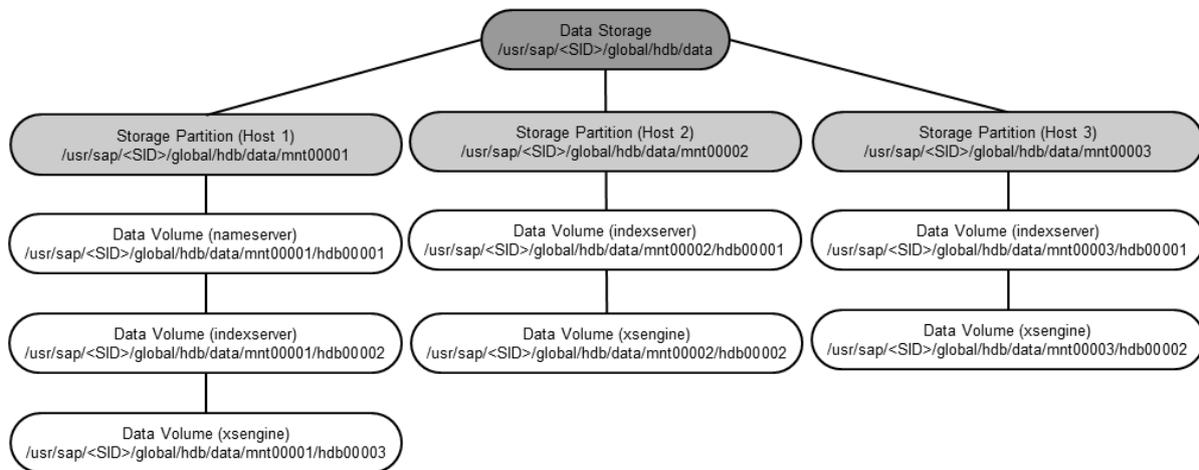
In a system with multitenant database containers, the volume names of tenant databases have a suffix to represent the database. For example, the indexserver volume for the first tenant database is `hdb00002.00002` and for the second database `hdb00002.000003`.

The services that persist data and therefore have volumes are the following:

Service	Note
nameserver	Only the nameserver service on the active master host persists data. Slave nameserver hosts communicate with the master, but do not persist data.
indexserver	The indexserver service on all hosts except standby hosts persists data.
xsengine (if running)	The xsengine service persists data on any host on which it is running.

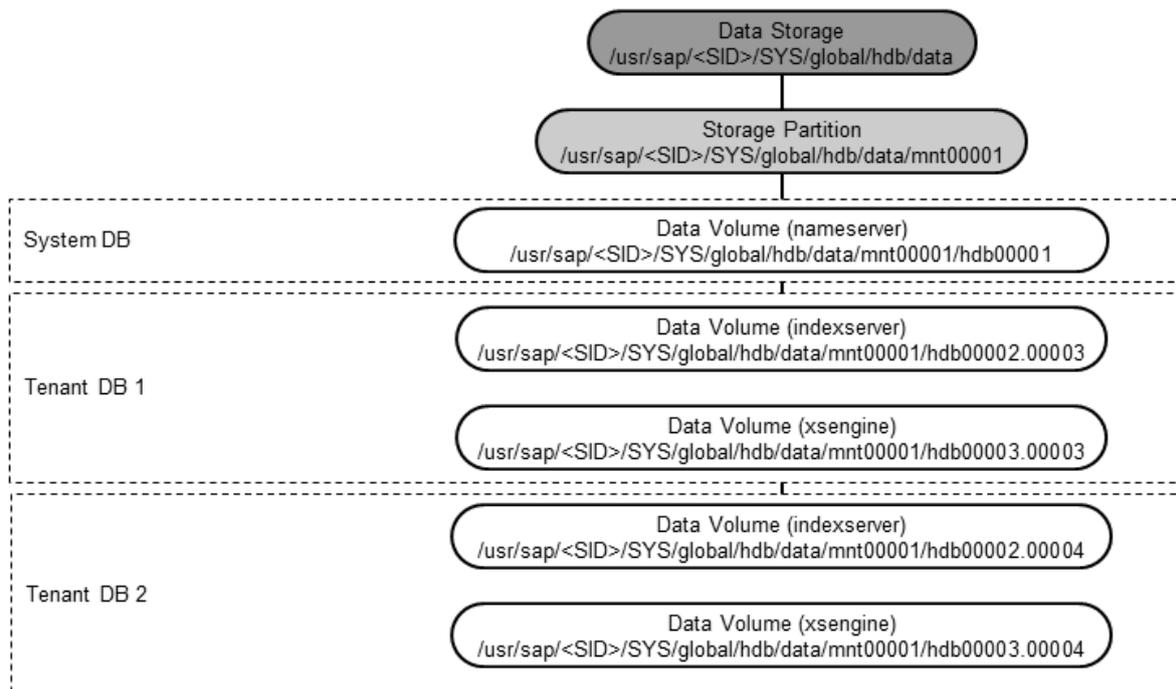
The following figure illustrates the default storage hierarchy described above using the example of data storage. The system is distributed across 3 hosts.

**Default Directory Hierarchy for Persistent Data Storage**



The following figure illustrates the default storage hierarchy described above using the example of data storage. The system is a system with multiple database containers on a single host.

**Directory Hierarchy for Persistent Data Storage (System with Multitenant Database Containers)**



## Data and Log Volumes

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). 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.

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.

### 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 2TB 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 *About Data Backups*.

### Related Information

[Data Backups \[page 881\]](#)

[Multitenant Database Containers \[page 15\]](#)

### 4.5.1.9.2 Hybrid LOBs (Large Objects)

With hybrid LOBs you can store LOB data on disk until needed rather than having it loaded into memory. This influences the size of the row store loaded into memory and therefore affects start up and takeover times.

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. When you import LOBs into the SAP HANA database you can configure, if they are loaded into memory or remain on disk. This capability is called hybrid LOBs. A hybrid LOB resides first on disk and is referenced only by an ID in the corresponding table column.

This significantly reduces main memory consumption, especially when LOB data is not actually requested. If there are memory shortages LOBs can be unloaded from main memory before column or table data needs to be unloaded. Thresholds are defined to keep only small LOBs (< 1000 bytes) in memory while larger LOBs are immediately transferred to disk and the reference is kept in table structures in memory.

With SPS 07 of SAP HANA new tables use hybrid LOBs instead of memory LOBs by default. For previous revision it is possible to alter tables using ALTER TABLE.

The parameter `memory threshold` can be adjusted in bytes for hybrid LOBs:

- >0 LOB data with length smaller or equal (<=) to `lob_memory_threshold` is stored in memory, LOBs bigger than `lob_memory_threshold` go on disk
- 0 all LOB data is stored on-disk (a VirtualFile is created per LOB)
- -1 all LOB data is stored in-memory (max supported length ~1GB) Default value: 1000 bytes

Example: CREATE TABLE

```
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
```

The default value for `memory threshold` is 1000, `memory threshold` is always referenced as smaller or equal (<=).

See the section Data Definition Statements in the *SAP HANA SQL and System Views Reference* for more information.

## Memory Consumption per Table

Even when a huge part of the 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](#)

### 4.5.1.9.2.1 Migrate to Hybrid LOBs

If you are using a revision prior to revision 70 you can manually migrate from disk based lobs to hybrid lobs.

#### Context

Migration is recommended for all revisions prior to SPS 07 (revision 70). Where this is not possible it is recommended to upgrade to a revision greater than or equal to revision 65. Then you should do a CSV export of the table and re-import it.

#### Procedure

1. Check for memory LOBs

```
SELECT
  SCHEMA_NAME,
  TABLE_NAME,
  COLUMN_NAME,
  MAP(CS_DATA_TYPE_NAME, 'ST_MEMORY_LOB', 'MEMORY', 'LOB', 'HYBRID')
LOB_STORAGE_TYPE,
MEMORY_THRESHOLD
FROM
  PUBLIC.TABLE_COLUMNS
WHERE
  DATA_TYPE_NAME IN ('BLOB', 'CLOB', 'NCLOB');
```

2. Switch to using hybrid LOBs

- a. To migrate memory LOBs to hybrid LOBs issue the following statements

```
ALTER TABLE t ALTER (data NCLOB MEMORY THRESHOLD 1000);
-- all LOBs <= 1000 bytes are in memory, larger LOBs are on disk
ALTER TABLE t ALTER (data CLOB ST_MEMORY_LOB);
-- explicit creation of (old) in-memory only LOBs
```

- b. After changing column store tables to hybrid LOB you should do a delta merge and LOB garbage collection.

```
merge delta of "SYSTEM"."T";
alter table "SYSTEM"."T" WITH PARAMETERS('LOB_GARBAGE_COLLECTION'='1');
```

---

## Results

You can check that your tables have been migrated to hybrid LOB by running the SQL statement from the first step again.

### 4.5.1.10 Monitoring Memory Usage

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 Note 1704499](#)

[Memory Indicators in the SAP HANA Studio \[page 282\]](#)

### 4.5.1.10.1 SAP HANA Used Memory

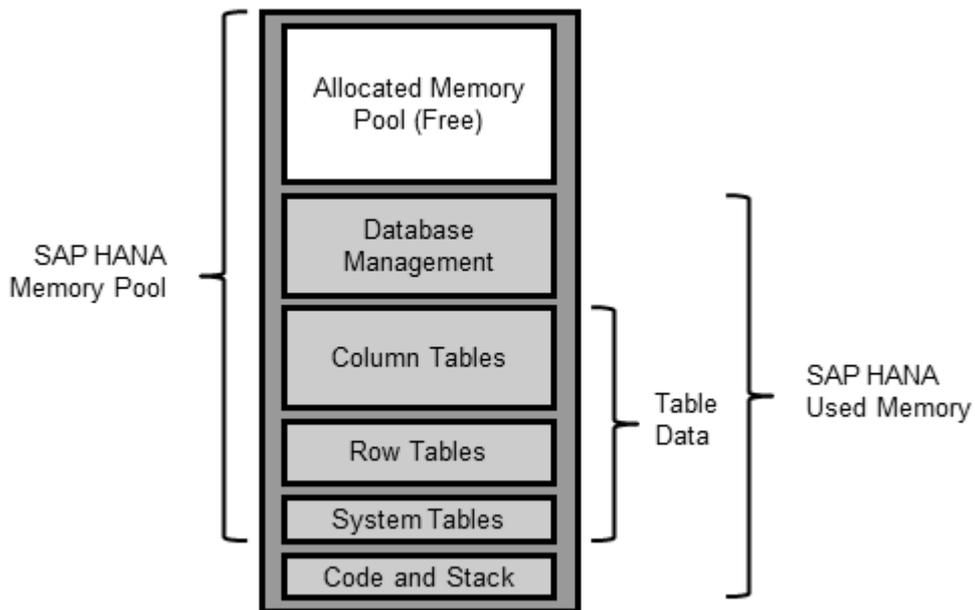
The total amount of memory used by SAP HANA is referred to as used memory. It includes program code and stack, all data and system tables, and the memory required for temporary computations.

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:

**SAP HANA Used Memory**



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**

[Reset Peak Used Memory \[page 277\]](#)

[Load/Unload a Column Table into/from Memory \[page 354\]](#)

[Setting a Memory Limit for SQL Statements \[page 277\]](#)

### **4.5.1.10.2 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](#) 

### 4.5.1.10.3 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 it 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.

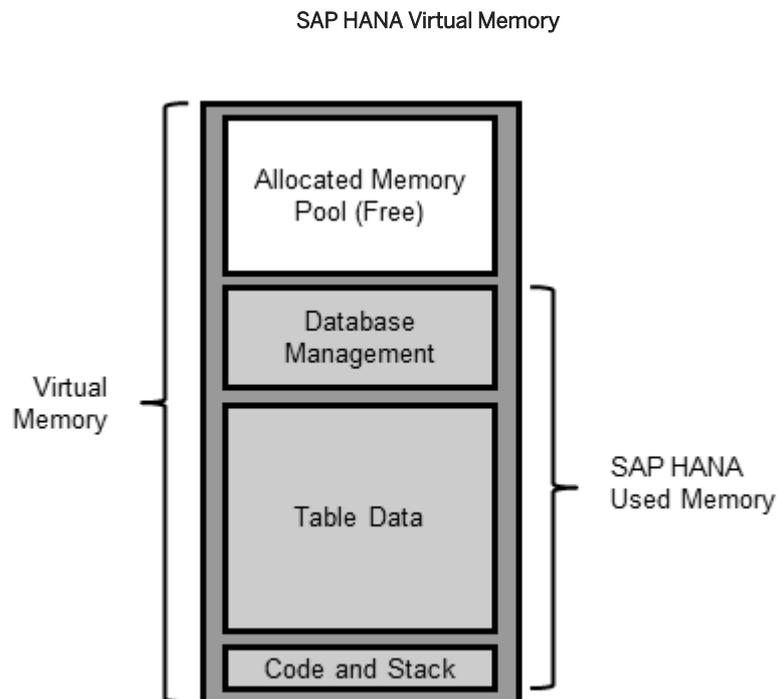
## Related Information

[Change the Global Memory Allocation Limit \[page 280\]](#)

### 4.5.1.10.4 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

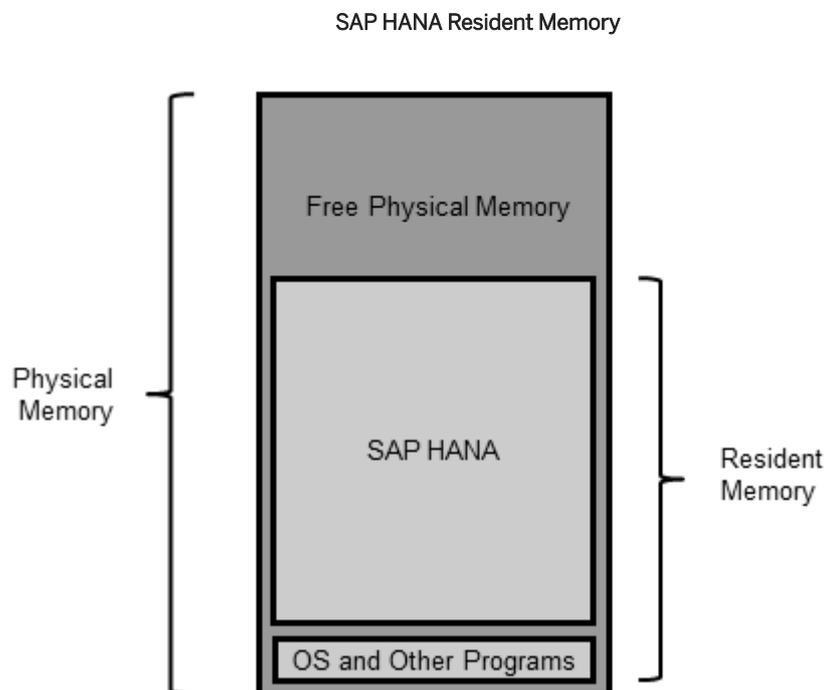
SAP HANA really consists of several separate processes, so the figure above shows all SAP HANA processes combined.

## 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:



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.

## **4.5.1.10.5 Reset Peak Used Memory**

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.

### **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.

### **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.

## **4.5.1.10.6 Setting a Memory Limit for SQL Statements**

The statement memory limit allows you to set a limit both per statement and per SAP HANA host.

### **Prerequisites**

To apply these settings you must have the system privilege INIFILE ADMIN.

## 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.

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.

This limit only applies to single statements, not the system as a whole. Tables which require much more memory than the limit applied here may be loaded into memory.

You can view the (peak) memory consumption of a statement in `M_EXPENSIVE_STATEMENTS.MEMORY_SIZE`.

Note that `M_EXPENSIVE_STATEMENTS.REUSED_MEMORY_SIZE` is not used as of SPS 09.

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.

## 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. In the `global.ini` file, expand the `memorymanager` section and set the parameter `statement_memory_limit`. Set a statement memory limit in GB (integer values only) with a value between 1 and some fraction of the global allocation limit.

### **i** Note

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).

Statements that exceed the limit you have set on a host are stopped by running out of memory.

3. In the `global.ini` file, expand the `memorymanager` section and set the parameter `statement_memory_limit_threshold` as a percentage of the global allocation limit (`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 the overall `global_allocation_limit` parameter setting.

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.

4. To set a user-specific statement limit and exclude a user from the global limit use the ALTER USER statement as shown here:

```
ALTER USER <user_name> SET PARAMETER STATEMENT MEMORY LIMIT = <gb>
```

- 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)

### **i** Note

Setting the statement memory limit to 0 will disable any statement memory limit for the user, or, to reset a user-specific limit use the CLEAR option:

```
ALTER USER <user_name> 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
<code>global_allocation_limit</code>	Default: 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
<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 GB > 69,12 GB (threshold is exceeded)

The statement is cancelled, as the threshold is exceeded, the `statement_memory_limit` is applied.

## Related Information

[Parameter Reference: Memory Consumption \[page 281\]](#)

### 4.5.1.10.7 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 in the `global.ini` configuration file.

## 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 value is the maximum allocation limit in MB. A missing entry or a value of 0 results in the system using the default settings. The global allocation limit is calculated by default 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.

## Procedure

1. In the Administration editor, choose the *Configuration* tab.  
The configuration files that contain the configuration information for the system are displayed.
2. Expand the `global.ini` configuration file and then the `memorymanager` section.
3. In the context menu for the `global_allocation_limit` parameter, choose *Change...*  
The *Change Configuration Value* dialog box appears.

4. Enter a value for the entire system and/or 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.

## Related Information

[Allocated Memory Pools and Allocation Limits \[page 273\]](#)

### 4.5.1.10.8 Parameter Reference: Memory Consumption

The `memorymanager` section of the `global.ini` file contains parameters that allow you to control the memory consumption of SAP HANA.

You can change the default settings in the configuration editor of the SAP HANA studio (recommended) or directly in the **global.ini** system properties file.

These parameters require tracking to be enabled in `global.ini` [`resource_tracking`]. Resource tracking creates a runtime overhead in SAP HANA:

- `enable_tracking = on`
- `memory_tracking = on`

#### **i** Note

In a system that supports multitenant database containers, you can configure the `global.ini` at both the system level and the database level. Parameters configured at the system level apply to the complete system and all databases. Parameters configured at the database level apply to the specified database only.

`global_allocation_limit` - limits the amount of memory that can be used by the system as a whole.

- The parameter `global_allocation_limit` defines the maximum memory allocation limit in MB.
- The global allocation limit is calculated by default 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. A missing entry or a value of 0 results in the system using the default settings.
- Does not require a restart. Available since SPS 08.

#### **i** Note

In a system that supports multitenant database containers, the global allocation limit configured at the system layer of the `global.ini` file is always effective regardless of any value configured at the database layer.

`statement_memory_limit` - defines the maximum memory allocation per statement in GB.

- 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.

- The default value is 0 (no limit). Set this parameter to a value between 1 GB and the value of the global allocation limit.
- Does not require a restart (applies to new statements). Available since SPS 09.

`statement_memory_limit_threshold` - defines the maximum memory allocation per statement as a percentage of the global allocation limit.

- If a value for this parameter has been set then the statement memory limit is only applied if the current SAP HANA memory consumption exceeds the statement memory limit threshold as a percentage of the global allocation limit.
- The default value is 0% (the `statement_memory_limit` is always respected). Set this parameter to a value between 1 GB and the value of the global allocation limit.
- Does not require a restart (applies to new statements). Available since SPS 09.

## Related Information

[Allocated Memory Pools and Allocation Limits \[page 273\]](#)

### 4.5.1.10.9 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	▮▮ <i>Landscape</i> ▸ <i>Services</i> ▮ 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	▮▮ <i>Landscape</i> ▸ <i>Services</i> ▮ 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	▮▮ <i>Landscape</i> ▸ <i>Services</i> ▮ tab of the <i>Administration</i> editor

## Related Information

[Monitoring Overall System Status and Resource Usage \[page 232\]](#)

[Monitoring Status and Resource Usage of System Components \[page 234\]](#)

## 4.5.1.11 Monitoring Using 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 on the various tab pages of the Administration editor. 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.

#### 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 system views in the following ways:

- On the *System Information* tab of the Administration editor  
Several predefined SQL SELECT statements on system views are listed here. 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*.

#### Tip

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, see *Use User-Defined SQL Statements for System Monitoring*.

- Query the view directly for example in the SQL console, or open it directly in the table editor  
You can search for a particular system view by right-clicking the *Catalog* entry in *Systems* view and choosing *Find Table*.

### 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.

The SAP HANA studio also includes 2 editors that allow you to visualize and explore historical resource utilization and memory allocation.

## System and Statistics Views in Multitenant Database Containers

Every multitenant database container system 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.

For more information, see *System and Statistics Views in Multiple-Container Systems*.

### Related Information

[Opening Tables and Views \[page 335\]](#)

[Use User-Defined SQL Statements for System Monitoring \[page 284\]](#)

[The Statistics Service \[page 245\]](#)

[Multitenant Database Containers \[page 15\]](#)

[System and Statistics Views in Multiple-Container Systems \[page 142\]](#)

## 4.5.1.12 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
<b>Create a new user-defined statement</b>	<ol style="list-style-type: none"><li>1. From the context menu, choose <i>New SQL Statement</i>.</li><li>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.</li><li>3. Save the statement.</li></ol>
<b>Import user-defined statements from file</b>	<ol style="list-style-type: none"><li>1. From the context menu, choose <i>Import SQL Statements</i>.</li><li>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.</li></ol>

### **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 \[page 286\]](#)

[Execute SQL Statements in SAP HANA Studio \[page 65\]](#)

### 4.5.1.12.1 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>
```

## 4.5.1.13 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 <a href="#">Overview</a> 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"><li>• <code>nameserver</code></li><li>• <code>indexserver</code></li><li>• <code>preprocessor</code></li></ul> <p>You can check the status of individual services of a system on the <a href="#">Landscape &gt; Services</a> tab.</p> <p>Normally, the <code>daemon</code> service automatically restarts inactive services, but you can also do so manually by choosing <a href="#">Restart Missing Services</a> from the context menu of the <a href="#">Landscape &gt; Services</a> tab.</p> <p>To investigate the reason for inactive services, consider the following actions:</p> <ul style="list-style-type: none"><li>• On the <a href="#">Diagnosis Files</a> 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.</li><li>• On the <a href="#">Diagnosis Files</a> tab, merge all diagnosis files and check the period before the service stopped (for example, the previous 30 minutes).</li><li>• Check any alerts generated in the period before the service stopped.</li></ul>

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 <a href="#">Overview</a> tab.</p> <div data-bbox="683 521 1396 680" style="background-color: #fff9c4; padding: 10px;"> <p><b>i Note</b></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 <a href="#">Alerts</a> 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 <a href="#">Overview</a> tab, and of individual services on the <a href="#">Landscape &gt; Services</a> tab.</p> <p>If memory usage bars are yellow or red, consider the following actions:</p> <ul style="list-style-type: none"> <li>• Check the error log for application or technical errors and contact SAP Support if necessary.</li> </ul> <div data-bbox="730 1211 1396 1384" style="background-color: #fff9c4; padding: 10px;"> <p><b>i Note</b></p> <p>To open the error log, from the main menu choose <a href="#">Window &gt; Show View &gt; General &gt; Error Log</a>.</p> </div> <ul style="list-style-type: none"> <li>• If there are no errors, increase available memory or reorganize your data.</li> </ul>

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 <a href="#">Overview</a> tab, and of individual services on the <a href="#">Landscape &gt; Services</a> tab.</p> <p>If the CPU usage bar is yellow or red, consider the following actions:</p> <ul style="list-style-type: none"> <li>• Check the error log for application or technical errors and contact SAP Support if necessary.</li> </ul> <div style="background-color: #fff9c4; padding: 10px; margin: 10px 0;"> <p><b>i Note</b></p> <p>To open the error log, from the main menu choose <a href="#">Window &gt; Show View &gt; General &gt; Error Log</a>.</p> </div> <ul style="list-style-type: none"> <li>• If there are no errors, analyze the trace files of running services. You can access these on the <a href="#">Diagnosis Files</a> tab.</li> <li>• Increase available resources.</li> </ul>
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 <a href="#">Overview</a> 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"> <li>• Reorganize your data.</li> <li>• Increase disk space.</li> </ul>
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 <a href="#">Overview</a> 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"> <li>• Reorganize unnecessary log files.</li> <li>• Verify that the last backup executed successfully.</li> </ul>
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 <a href="#">Overview</a> 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"> <li>• Switch off traces.</li> <li>• Delete unnecessary trace files.</li> <li>• Review the configured trace file rotation.</li> </ul>

Step	What to Check	What to Do
8	Regular, successful execution of backups	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.  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.
9	Sufficient space for backups	Ensure that there is sufficient space at the chosen backup destination.
10	Existence of crash dump files	You can check for crash dump files on the <i>Diagnosis Files</i> tab.  If such files exist, investigate further. If necessary, contact SAP Support.
11	Number of active threads and the duration of the top 5 threads	You can review active threads on the <b>► Performance ► Threads ▾</b> tab.  If further investigation is required, refer to the other sub-tabs of the <i>Performance</i> tab. Here you can analyze the following: <ul style="list-style-type: none"> <li>• Expensive SQL statements</li> <li>• Sessions</li> <li>• SQL performance history</li> <li>• Progress of long-running operations</li> <li>• System load history</li> </ul>
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 \[page 232\]](#)

[Monitoring Alerts \[page 242\]](#)

[View Diagnosis Files in SAP HANA Studio \[page 461\]](#)

[Monitoring Memory Usage \[page 270\]](#)

[Monitoring Status and Resource Usage of System Components \[page 234\]](#)

[Monitoring Disk Space \[page 261\]](#)

[Configure Traces in SAP HANA Studio \[page 465\]](#)

[Configure Trace File Rotation \[page 480\]](#)

[Monitoring System Performance \[page 253\]](#)

## 4.5.2 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 allocation history 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.

For more information about monitoring tenant databases, see *Managing Multitenant Database Containers*.

#### **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

[SAP HANA Cockpit \[page 22\]](#)

[Open SAP HANA Cockpit \[page 23\]](#)

[Managing Multitenant Database Containers \[page 104\]](#)

[Monitor and Analyze Past Performance \[page 313\]](#)

[Monitoring Alerts \[page 300\]](#)

[Configuring Alerts \[page 304\]](#)

### 4.5.2.1 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, for more detailed monitoring and analysis, drill down into the dedicated apps.

#### **i** Note

If you removed any of the tiles from the homepage of the SAP HANA cockpit, you can add them again from the tile catalog. For more information, see *Customizing the Homepage of SAP HANA Cockpit*.

## Related Information

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

### 4.5.2.1.1 Monitor Status and Resource Usage of Database Services

To monitor the health of your SAP HANA database in more detail, for example, to troubleshoot performance bottlenecks, you analyze the status and resource usage of individual database services. If necessary, you can perform follow-up operations, for example, start missing services, or stop or kill a service.

#### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Manage Services](#) tile is visible on the homepage of the SAP HANA cockpit. This tile is visible by default, but if you removed it, you can add it again from the tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

#### Procedure

Open the [Manage Services](#) app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.

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 in the [Alerts](#) app by clicking [Go to Alerts](#) in the footer bar.
- If you want to investigate the memory usage history of a particular service, click the value in the [Used Memory](#) column to open the [Memory Allocation Statistics](#) app for the service in a new window.

- 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*.

## Related Information

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Service Details \[page 293\]](#)

[Operations on Services \[page 295\]](#)

[Monitoring Alerts \[page 300\]](#)

[Analyze Memory Allocation Statistics \[page 297\]](#)

[Assign Roles to a User \[page 717\]](#)

### 4.5.2.1.1.1 Service Details

The *Monitor Services* app provides you with detailed information about database services.

#### **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"> <li>• <i>Running</i></li> <li>• <i>Not Running</i></li> </ul> <p>To investigate why the service is not running, you can navigate to the crash dump file created when the service stopped.</p> <div data-bbox="560 1727 1398 1944" style="background-color: #fff9c4; padding: 10px; margin-top: 10px;"> <p><b>i</b> 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

Column	Description
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"> <li>• <i>Master</i> The service is the active master worker.</li> <li>• No entry The service is a slave worker.</li> <li>• <i>Standby</i> The service is in standby mode. It does not contain any data and does not receive any requests.</li> </ul>
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: #fff9c4; padding: 5px;"> <p><b>i Note</b></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>
Memory	<p>Mini chart visualizing the memory usage of the service</p> <p>Clicking the mini chart opens the <i>Memory Allocation Statistics</i> 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 <i>Memory Allocation Statistics</i> 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

Column	Description
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

[Monitoring Memory Usage \[page 270\]](#)

[Analyze Memory Allocation Statistics \[page 297\]](#)

### 4.5.2.1.1.2 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 the *Monitor Services* app. 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 the `sap.hana.admin.roles::Administrator` role. Also, depending on the service, some options may not be available.

Option	Description
<i>Start Missing Services</i>	Starts any inactive services
<i>Stop</i>	Stops the selected service normally The service is then typically restarted.

Option	Description
<i>Kill</i>	<p>Stops the selected service immediately and if the related option selected, creates a crash dump file</p> <p>The services is then typically restarted.</p> <p><b>i Note</b></p> <p>You can access the generated crash dump file in the SAP HANA studio. For more information, see <i>Diagnosis Files</i> in the <i>SAP HANA Administration Guide</i>.</p>
<i>Remove</i>	<p>Removes the selected service</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"> <li>• Name server</li> <li>• Master index server</li> <li>• Primary index server on a host</li> </ul> <p><b>i Note</b></p> <p>To remove a service, you must have the EXECUTE privilege on the stored procedure SYS.UPDATE_LANDSCAPE_CONFIGURATION.</p>
<i>Reset Memory Statistics</i>	<p>Resets all memory statistics for all services</p> <p>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. 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> <p>For more information about memory concepts in SAP HANA, see the section on memory usage.</p>

### **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

[Monitoring Memory Usage \[page 270\]](#)

[High Availability for SAP HANA \[page 774\]](#)

## 4.5.2.1.2 Analyze Memory Allocation Statistics

Analyzing the memory allocation history of the SAP HANA database can help you to investigate out-of-memory situations, memory corruptions, memory leaks and so on.

### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Manage Services](#) tile is visible on the homepage of the SAP HANA cockpit. This tile is visible by default, but if you removed it, you can add it again from the tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

### Context

The [Memory Allocation Statistics](#) app enables you to visualize and explore the memory allocation of every service. The following system views provide the information with which values for current and historical memory allocation are calculated:

- M\_SERVICE\_MEMORY (SYS)
- M\_HEAP\_MEMORY (SYS)
- HOST\_SERVICE\_COMPONENT\_MEMORY (\_SYS\_STATISTICS)
- HOST\_HEAP\_ALLOCATORS (\_SYS\_STATISTICS)

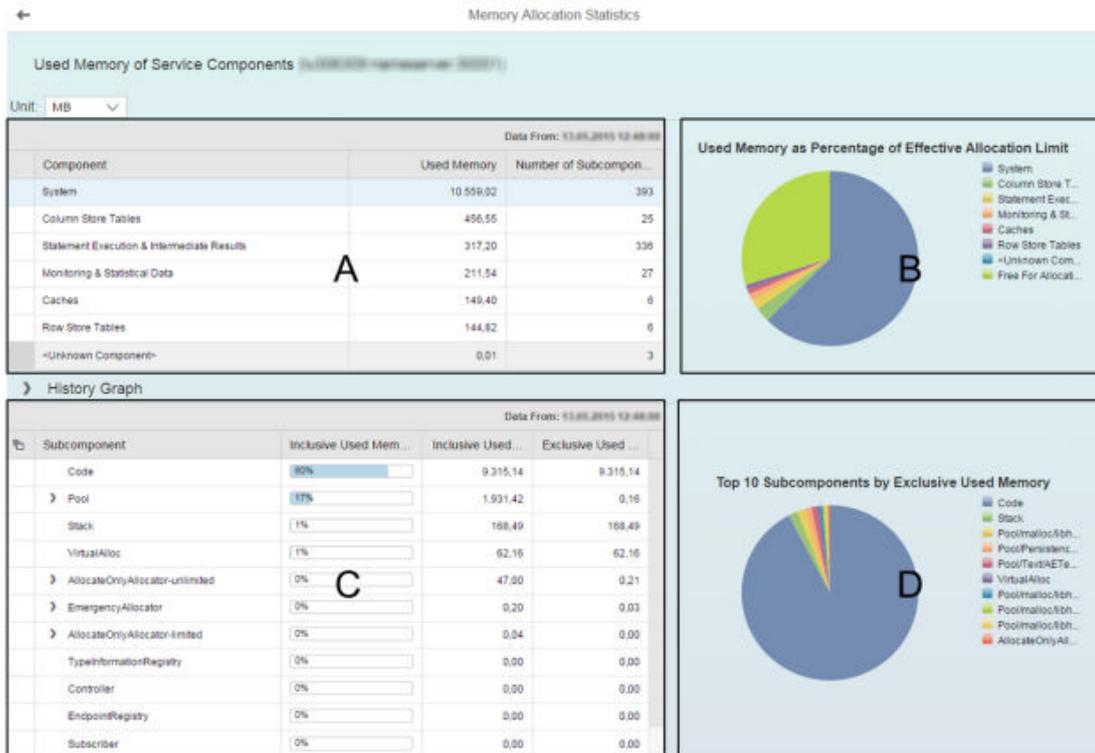
For more information about these views, see the *SAP HANA SQL and System Views Reference*.

### Procedure

1. Open the [Manage Services](#) app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.  
You see the status of all the services in the system. For each service, detailed information about its memory consumption is available. For more information, see [Service Details](#).
2. Open the [Memory Allocation Statistics](#) app by clicking the value in [Used Memory](#) column of the service whose memory allocation history you want to see.

#### → Tip

If the [Used Memory](#) column is not visible, click the configuration button in the footer toolbar and select it.



Initial View of Memory Allocation Statistics App

The following information is displayed in screen areas identified above:

- A: The components of the selected service listed in descending order of current used memory (default)
  - B: Current breakdown of the component's used memory displayed as a pie chart
  - C: Subcomponents of the selected component listed in descending order of current used inclusive memory (default)
  - D: Current breakdown of memory usage of the 10 highest consuming subcomponents displayed as a pie chart
3. Analyze the used memory history of the component and its allocators by opening the history graph and exploring the data.

Several options are available, for example:

- To change the visualized time period, adjust the *From/To* values.
- To visualize the memory usage of individual allocators over the same time period as the selected component, simply select them in the allocators table.
- To see time-specific values, click a point on the graph or hover over with the mouse.

## Related Information

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Monitoring Memory Usage \[page 270\]](#)

## 4.5.2.1.3 Monitor Tables

Monitor tables to optimize resource utilization and improve query performance.

### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`.
- The *User Tables* tile is visible on the homepage of the SAP HANA cockpit by default. If it is not visible, you can add it from the tile catalog. For more information, see *Customizing the Homepage of SAP HANA Cockpit*.

### Context

The *User Tables* app shows the comparative memory utilization by host and the number of high and medium priority alerts.

### Procedure

1. Open the *User Tables* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit. 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 *Search*. 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  (*Show table history as table*).
- For graphical format display, choose  (*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.

### Caution

Be aware that you need additional licenses for SAP HANA options such as SAP HANA dynamic tiering and smart data streaming. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities*.

## Related Information

[Alert Details \[page 302\]](#)

[Alert Priorities \[page 243\]](#)

[Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities \[page 1360\]](#)

[SAP Note 2092669 - Release Note SAP HANA Data Warehousing Foundation](#)

## 4.5.2.2 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).

### Note

Alert monitoring and configuration is possible with the SAP HANA cockpit only if the monitoring and alerting functions in the system are being implemented by the **embedded statistics service**, not the statistics server. For more information about migrating to the statistics service, see SAP Note 1917938.

---

## Related Information

[Configuring Alerts \[page 304\]](#)

[The Statistics Service \[page 245\]](#)

[Monitor Alerts \[page 301\]](#)

[SAP Note 1917938](#)

### 4.5.2.2.1 Monitor Alerts

Regularly monitoring alerts ensures that you can take timely and appropriate action in the event of a problem with SAP HANA.

#### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Alerts](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the [SAP HANA Database Administration](#) tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

#### Procedure

1. Open the SAP HANA cockpit.  
The number of high and medium priority alerts in the database is displayed on the [Alerts](#) tile.  
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 [Search](#) field (for example, **license**).
  - To filter according to additional attributes including priority and date of occurrence, click on the filter icon in the footer bar and select the required filter(s).
2. Open the [Alerts](#) app by clicking the tile of the same name.  
All high and medium priority alerts are displayed in list format on the left. To see more detailed information about a specific alert on the right, simply select it.  
  
To see all past alerts, click [Past Alerts](#) in the footer toolbar. You can also use the options available for filtering, searching, and sorting to customize the list of alerts. For example, if you wanted to see all memory-related alerts in last week, you could filter by check category and time frame.

## Related Information

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Alert Details \[page 302\]](#)

[Alert Priorities \[page 243\]](#)

### 4.5.2.2.1.1 Alert Details

When you select an alert in the *Alerts* app, 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 Related Information.
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 \[page 808\]](#)

## 4.5.2.2.1.2 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

## 4.5.2.2.2 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.

### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Alerts](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the [SAP HANA Database Administration](#) tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

### 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 \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Assign Roles to a User \[page 717\]](#)

### 4.5.2.2.3 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

#### Note

You can configure alerts with the SAP HANA cockpit only if the monitoring and alerting functions in the system are being implemented by the **embedded statistics service**, not the statistics server. For more information about migrating to the statistics service, see SAP Note 1917938.

## Related Information

[Configure Alerting Thresholds \[page 306\]](#)

[Switch Alerting Off/On \[page 308\]](#)

[Set Up E-Mail Notification \[page 310\]](#)

[Check for Alerts Out of Schedule \[page 312\]](#)

[SAP Note 1917938 !\[\]\(cba1c7962ef6706b09e0f24cde21949d\_img.jpg\)](#)

## 4.5.2.2.3.1 Alert Checker Details

When you select an alert checker in the *Alert Configuration* app, 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.  <b>i Note</b> 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 <b>cannot</b> 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 \[page 306\]](#)

[Configure Alerting Thresholds \[page 306\]](#)

## 4.5.2.2.3.2 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"><li>• INTERVALLENGTH</li><li>• SKIP_INTERVAL_ON_DISABLE</li></ul> <p>Once <math>INTERVALLENGTH \times SKIP\_INTERVAL\_ON\_DISABLE</math> 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 the <a href="#">Alert Configuration</a> app.</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 \[page 308\]](#)

## 4.5.2.2.3.3 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.

### Prerequisites

- You have the privileges granted by `sap.hana.admin.roles::Administrator`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.

- The [Configure Alerts](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the [SAP HANA Database Administration](#) tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

## 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 the [Alert Configuration](#) app by clicking the [Configure Alerts](#) tile on the homepage of the SAP HANA cockpit.

### Note

It is also possible to navigate to the [Alerts Configuration](#) app from the [Alerts](#) app by clicking [Configure Alert](#) in the footer toolbar.

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 \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Alert Checker Details \[page 305\]](#)

[Assign Roles to a User \[page 717\]](#)

### 4.5.2.2.3.4 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.

## Prerequisites

- You have the privileges granted by `sap.hana.admin.roles::Administrator`. You can grant roles using the *Assign Roles* app of the SAP HANA cockpit. For more information, see *Assign Roles to a User* in the *SAP HANA Administration Guide*.
- The *Configure Alerts* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA Database Administration* tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

## 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.

### Note

The system automatically switches failed alert checkers back on after a certain length of time. For more information, see *Alert Checker Statuses*.

## Procedure

1. Open the *Alert Configuration* app by clicking the *Configure Alerts* tile on the homepage of the SAP HANA cockpit.

### **i** Note

It is also possible to navigate to the *Alerts Configuration* app from the *Alerts* app by clicking *Configure Alert* in the footer toolbar.

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 \[page 23\]](#)

[Alert Checker Details \[page 305\]](#)

[Alert Checker Statuses \[page 306\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

## 4.5.2.2.3.5 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.

### Prerequisites

- You have the privileges granted by `sap.hana.admin.roles::Administrator`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Configure Alerts](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA Database Administration* tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

### Context

If you want to be notified by e-mail about new alerts when they are issued, you can set this up in the [Alerts Configuration](#) app. 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 the [Alert Configuration](#) app by clicking the [Configure Alerts](#) tile on the homepage of the SAP HANA cockpit.

#### **i** Note

It is also possible to navigate to the [Alerts Configuration](#) app from the [Alerts](#) app by clicking [Configure Alert](#) in the footer toolbar.

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 (see step 4).

- In the footer toolbar, click the envelope icon and choose *Default Recipient(s)*.
- Enter the e-mail addresses of the recipients.
- 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.

- Find the alert checker that you want to configure.  
The detailed configuration of the alert checker is displayed on the right.
- Open the alert checker for editing by clicking *Edit*.
- In the *Email* field, enter the e-mail addresses of the recipients.
- 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 \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Alert Checker Details \[page 305\]](#)

[Alert Checker Statuses \[page 306\]](#)

## 4.5.2.2.3.6 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.

### Prerequisites

- You have the privileges granted by `sap.hana.admin.roles::Administrator`. You can grant roles using the *Assign Roles* app of the SAP HANA cockpit. For more information, see *Assign Roles to a User* in the *SAP HANA Administration Guide*.
- The *Configure Alerts* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA Database Administration* tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

### 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 the *Alert Configuration* app by clicking the *Configure Alerts* tile on the homepage of the SAP HANA cockpit.

#### **i** Note

It is also possible to navigate to the *Alerts Configuration* app from the *Alerts* app by clicking *Configure Alert* in the footer toolbar.

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 \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Alert Checker Details \[page 305\]](#)

[Alert Checker Statuses \[page 306\]](#)

[Switch Alerting Off/On \[page 308\]](#)

### 4.5.2.3 Monitoring Performance

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.

#### 4.5.2.3.1 Monitor and Analyze Past Performance

Analyzing the performance of the SAP HANA database over time can help you to pinpoint bottlenecks, identify patterns, and forecast requirements. Use the *Performance Monitor* app to visually analyze historical performance data across a range of key performance indicators related in particular to memory, disk, and CPU usage.

#### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`. You can grant roles using the *Assign Roles* app of the SAP HANA cockpit. For more information, see *Assign Roles to a User* in the *SAP HANA Administration Guide*.
- The *Used Memory*, *Disk Usage*, or *CPU Usage* tiles are visible on the homepage of the SAP HANA cockpit. If they're not, you can add them again from the tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

#### Procedure

Open the *Performance Monitor* app by clicking the *CPU Usage*, *Disk Usage*, or *Used Memory* tile on the homepage of the SAP HANA cockpit.

The *Performance Monitor* app opens displaying the load graph for the selected resource: CPU, disk, or memory. The load graph initially visualizes resource usage on the master host and master index or name server according to the default KPI group of the selected resource.

You can customize the information displayed on the load graph in several ways, for example:

- Show additional KPIs and create customized KPI group variants  
For a list of all available KPIs, see *Key Performance Indicators*.
- Add additional hosts and services
- Show which jobs had an effect on system performance.  
Select *Show Jobs* to display jobs above the load graph.
- Increase or decrease the date range of collected data
- Zoom in to a specific time

## Related Information

[Key Performance Indicators \[page 314\]](#)

### 4.5.2.3.1.1 Key Performance Indicators

The *Performance Monitor* app allows you select a range of host- and service-level KPIs to analyze historical performance data of the SAP HANA database.

#### Host KPIs

KPI	Description
CPU	CPU used by all processes relative 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 indexserver 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
Active transaction ID range	Range between newest and oldest active transaction 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

KPI	Description
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

## Related Information

[Monitoring Memory Usage \[page 270\]](#)

### 4.5.2.3.2 Analyze Critical Statements

Analyzing the current most critical statements running in the 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.

## Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Monitor Statements](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it again from the tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

## Context

The [Monitor Statements](#) app lists and provides you with information about the 100 most critical statements 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

## Procedure

1. Open the *Monitor Statements* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.

The 100 most critical statements currently active in the database are listed. By default, statements are listed in order of longest runtime. For each statement, you can see the full statement string, as well as the ID of the session in which the statement is running. You can identify the application, the application user and the database user running the statement, and whether the statement is related to a blocking transaction.

2. Optional: Monitor the memory consumption of statements by clicking *Enable Memory Tracking* in the footer toolbar.

Information about the memory consumption of statement execution 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. Optional: If a statement is involved in a blocked transaction or using an excessive amount of memory, cancel the session the statement is running in (or the blocking session) by clicking *Cancel Session* in the footer toolbar.

## Related Information

[Statement Details \[page 317\]](#)

[Setting a Memory Limit for SQL Statements \[page 277\]](#)

### 4.5.2.3.2.1 Statement Details

The *Monitor Statements* app provides you with detailed information about the 100 most critical statements running 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.

## General Statement Information

The table below lists the information available for statements.

Detail	Description
Statement Runtime	Runtime of the current execution of the statement
SQL Statement	The full SQL string
Session	ID of the session in which the statement is running
Lock Wait Time	Lock wait time of the statement if blocked
Cursor Duration	Statement execution time including communication time with clients
Application	Application from which the statement was executed
Application User	Application user who executed the statement
Database User	Database user who executed the statement
Blocking Session	ID of the session that is blocking execution of the statement

## Statement Memory Information

If statement memory tracking is enabled, the following additional information for each statement is available:

Detail	Description
Used Memory (MB)	Current snapshot of memory used to execute the statement
Allocated Memory (MB)	Current snapshot of amount of memory allocated to execute the statement
Average Execution Memory (MB)	Average amount of memory used to execute the statement
Maximum Execution Memory (MB)	Maximum amount of memory used to execute the statement
Minimum Execution Memory (MB)	Minimum amount of memory used to execute the statement
Total Execution Memory (MB)	Total amount of memory used to execute the statement
Workload Class	Workload class  The workload class influences dynamic resource consumption at the session or statement level. For more information, see <i>Managing Workload</i> in the <i>SAP HANA Administration Guide</i> .

## Related Information

[Managing Workload with Workload Classes \[page 435\]](#)

---

### 4.5.2.3.3 Analyze Threads

You can monitor the longest-running threads active in your system. 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.

#### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see *Assign Roles to a User* in the *SAP HANA Administration Guide*.
- The *Threads* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it again from the tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

#### Context

The *Threads* app lists and provides you with information about the 1000 longest-running threads currently active in the database.

#### Procedure

1. Open the *Threads* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.  
The 1000 longest-running threads currently active in the database are listed. By default, threads are listed in order of longest runtime. For each statement, you can see the duration, as well as the name of the service which is executing the thread. You can identify the host, the post and the thread type, and whether the statement is related to a blocking transaction.
2. Optional: If a thread is involved in a blocked transaction or using an excessive amount of memory, cancel the operation executing the thread by clicking [Cancel Operations](#) in the footer toolbar.

#### Related Information

[Thread Details \[page 320\]](#)

## 4.5.2.3.3.1 Thread Details

The *Threads* app 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)
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

Detail	Description
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
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

Detail	Description
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
Waiting Schema Name	Name of the schema on which the lock is currently placed

#### 4.5.2.3.4 Capturing and Replaying Workloads

Capturing and replaying workloads from an SAP HANA database can help you evaluate potential impacts on performance or stability after a change in hardware or software configuration.

The following sections provide an overview of the capture and replay tool:

#### What is SAP HANA Capture and Replay?

SAP HANA capture and replay is a tool that allows you to capture the workload of a productive system and to replay the captured workload on a target system. Furthermore, the tool allows you to analyze runtimes and compare performance between these different systems.

##### **i** Note

If you are using Revision 1.00.122.14 or higher, refer to the SAP HANA 2.0 SPS02 version of the documentation. For more information, see *SAP Note 2362820*.

#### What is a workload?

Workload is defined as the amount of work to be produced on SAP HANA database systems in a specified period of time. In the context of SAP HANA capture and replay, workload can mean any change to the database via SQL (for example, tables, views), analytical or transactional queries to the database via SQL, as well as internal tasks and housekeeping operations.

The workload can be created by applications or clients (for example, NetWeaver worker processes, BOBJ Frontends).

### **i** Note

When using a Revision of SAP HANA Platform 1.0 below Revision 122.14, the tool cannot capture and replay workload coming from native XS Classic applications.

## Sample use cases

Possible use cases are:

- Hardware change
- SAP HANA revision upgrade
- SAP HANA ini file change
- Table partitioning change
- Index change
- Landscape reorganization for SAP HANA scale-out systems

## How does SAP HANA Capture and Replay work?

The main steps involved in the capturing and replaying process are:

### 1. Capture

You can capture the entire workload of a productive system. In this step the tool collects the execution context information automatically together with the incoming requests to the database. In addition, workload capture stores the start times of statements, as well as user specified fetch size and parameter sets.

### **i** Note

Copy manually the captured file and the database backup from the production system to the target system.

### 2. Preprocess

In this step the tool reconstructs and optimizes the captured workload files to make them replayable on a target system. This process is a one-time operation and the stored preprocessed files are reusable for multiple replays.

### 3. Replay and Analyze

In this step the tool replays the preprocessed workload in the correct order based on the statement timestamp. Together with the collected execution context, it allows you to accurately simulate the database workloads. Furthermore, it allows you to generate a summary of the source workload, as well as the result of the replay compared to that source capture.

There are two types of workload reports:

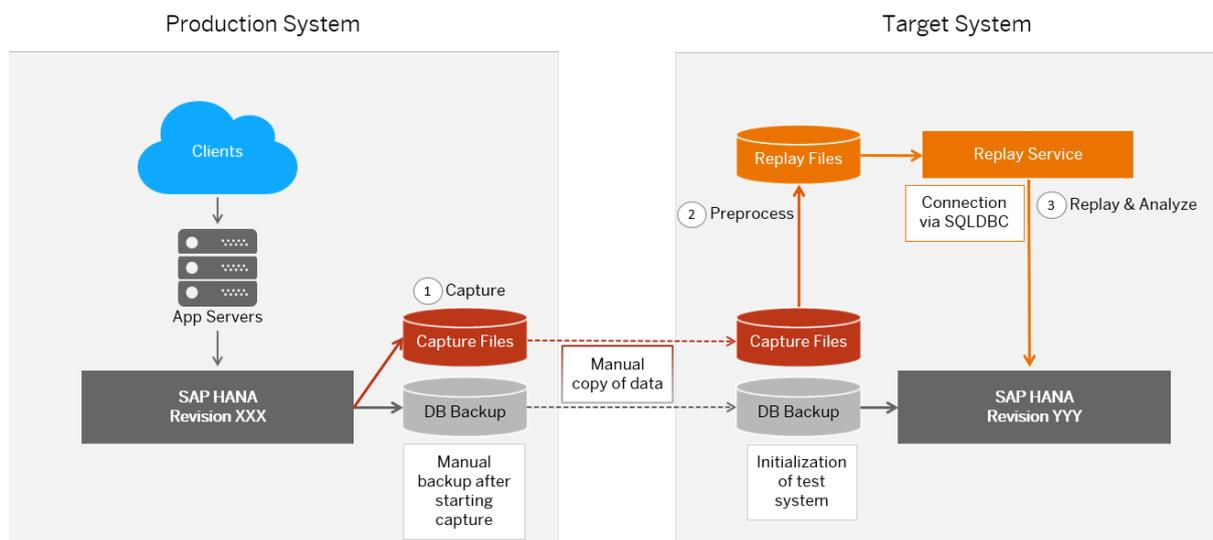
- The report of the source system, which is generated from the captured workload data.

- The report of the difference between the source system and the replayed system. This report includes system utilization (for example, CPU, Wait, and Throughput) and query runtime.

The following interactive graphic provides you with 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 \[page 325\]](#)
- [Preprocess a Captured Workload \[page 326\]](#)
- [Replay a Preprocessed Workload \[page 327\]](#)

## Related Information

[Capture a Workload \[page 325\]](#)

[Preprocess a Captured Workload \[page 326\]](#)

[Replay a Preprocessed Workload \[page 327\]](#)

[SAP Note 2362820](#)

## 4.5.2.3.4.1 Capture a Workload

You can capture and monitor the workload from an SAP HANA system.

### Prerequisites

- Install the *HANA WORKLOAD REPLAY 1.0* delivery unit. For more information, see *Deploy a Delivery Unit Archive (\*.tgz)*.
- You have the privileges granted by the role `sap.hana.replay.roles::Capture`. You can assign roles using the *Assign Roles* app of the SAP HANA cockpit. For more information, see *Assign Roles to a User* in the *SAP HANA Administration Guide*.
- The *Capture Workload* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it again from the tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

#### ➔ Recommendation

To ensure that the target system and the source system are in a consistent state for the capture and replay process, it is recommended to perform a full database backup after starting the capture.

### Context

The *Capture Workload* app lists and provides you with information about the captured workloads.

### Procedure

1. Open the *Capture Workload* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.
2. In the *Capture Management* display area, click *Start New Capture* on the bottom right.
3. Configure the workload and click *Start Capture* on the bottom right.

The *Capture Monitor* opens displaying information about the configured capture.

#### **i** Note

The captured file is stored under the tracedir with an \*.cpt file extension.

4. Optional: Review the captured workload by clicking the workload with the status *Captured* in the *Capture Management* display area.

The *Capture Report* display area provides information about the captured workload.

## Related Information

[Deploy a Delivery Unit Archive \(\\*.tgz\) \[page 632\]](#)

[Assign Roles to a User \[page 717\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

### 4.5.2.3.4.2 Preprocess a Captured Workload

The preprocessing step is required to optimize captured files before replaying them.

#### Prerequisites

- You have the privileges granted by the role `sap.hana.replay.roles::Replay`. You can assign roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Replay Workload](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it again from the tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).
- You have captured workloads using the [Capture Workload](#) app. For more information, see [Capture a Workload](#).
- Copy the captured files from the production system to the test system in the `tracedir` directory.

#### **i** Note

It is highly recommended to perform the preprocessing in the target system or in a separate control system, not in the production system. The preprocessing may require significant computing power.

#### Context

The [Replay Workload](#) app lists and provides you with information about replayed workloads, workloads ready for replay, as well as workloads being preprocessed.

#### Procedure

1. Open the [Replay Workload](#) app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.

The captured workloads are listed in the [Replay Management](#) display area. By default, workloads are listed in order of the last captured workload.

2. Click [Edit](#) and check the captured workload that you want to preprocess.
3. Click [Start Preprocessing](#) on the bottom right.
4. Optional: Open the [Replay Candidate Details](#) display area by clicking a workload with the status [Preprocess Completed](#).
5. Optional: Start the replay from the [Replay Candidate Details](#) display area by clicking [Configure Replay](#).

## Related Information

[Assign Roles to a User \[page 717\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Capture a Workload \[page 325\]](#)

### 4.5.2.3.4.3 Replay a Preprocessed Workload

You can replay all preprocessed workloads as often as necessary.

## Prerequisites

- You have the privileges granted by the role `sap.hana.replay.roles::Replay`. You can assign roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Replay Workload](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it again from the tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).
- You have preprocessed the captured workloads using the [Replay Workload](#) app. For more information, see [Preprocess a Captured Workload](#).
- Before starting to replay a preprocessed workload keep also a `.cpt` file copy in the `tracedir` location of the target system.
- The replayer service is running. For more information, see [Start the Replayer Service](#).
- If a database backup is available, restore the database before starting the replay.

## Procedure

1. Open the [Replay Workload](#) app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.

The captured workloads are listed in the [Replay Management](#) display area. By default, workloads are listed in order of the last preprocessed workload.

2. Open a workload with the status [Preprocess Completed](#) by clicking on it.

3. Click [Configure Replay](#) in the [Replay Candidate Details](#) display area.

Enter the required information for the following areas:

- In the [Target Instance Information](#) area enter the host name, the instance number, and select the database mode.
- In the [Replayer Options](#) area enter the replay admin user, the replay admin password and select the replay speed.
- In the [User Authentication](#) area enter the user name and password/secure store key.
- In the [Replayer List](#) area select the desired replayer service that should be used for the replay.

4. Click [Start Replay](#).

The [Replay Management](#) opens displaying the replayed workloads in the [Replay List](#) tab.

5. Open the [Replay Monitor](#) by clicking one of the workloads displayed in the [Replay Management](#).
6. Monitor the replayed workload by clicking [Go to Report](#) on bottom right after the replay is completed.

The [Replay Report](#) opens displaying replay and capture configuration details in the [Summary](#) tab. The [Compare](#) tab displays a comparison of the SQL statements by execution time.

7. Optional: To find more information on an SQL statement, open the [Performance Comparison Detail](#) by clicking the corresponding row.

## Related Information

[Assign Roles to a User \[page 717\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Preprocess a Captured Workload \[page 326\]](#)

[Start the Replayer Service \[page 328\]](#)

### 4.5.2.3.4.3.1 Start the Replayer Service

This procedure documents the necessary steps to start the replayer service, which is needed for replaying workloads. The replayer service reads SQL commands from the preprocessed capture and executes them one-by-one in timestamp-based order.

## Prerequisites

You have the WORKLOAD REPLAY ADMIN system privilege.

## Procedure

1. Create a new configuration file as <code><\$SAP\_RETRIEVAL\_PATH/wlreplayer.ini>.
2. Fill in the following information in the wlreplayer.ini

```
[communication]
listeninterface = .global
```

```
[trace]
filename = wlreplayer
alertfilename = wlreplayer_alert
```

Wlreplayer is not a daemon, but a process, which needs to be ended when the replayer service is not needed anymore.

3. Start the replayer service using the following command:

```
hdbwlreplayer -controlhost <controlHost> -controlinstnum <controlInstanceNumber>
-controladminkey <userName,secureStoreKey> [-controldbname
<controlDatabaseName>] -port <listenPortNumber>
```

### **i** Note

Running the command over the target system does not trigger the replay, it only starts the replayer service.

#### Parameter Description

Parameter	Description
controlhost	Controls the database host name for replayer service (without a sqlport)
controlinstnum	Controls the database instance number
controladminkey	User name and secure store key of control management connection separated by a comma
controldbname	Controls the database name in case of multitenant database containers
port	Discretionary port number for internal communication

4. Optional: Check if the replayer is associated with the control database using the following query:

```
CALL WORKLOAD_REPLAY('get_replayer_list', NULL)
```

#### Result:

```
"replayer_list":[
  {
    "capture_id":0,
    "replay_id":0,
    "replayer_host":"selqnta4",
    "replayer_id":33682,
    "status":"IDLE"
  }
]
```

```
}
```

5. To stop the process use Ctrl+C.

## 4.5.2.3.5 Analyzing Workloads

Analyzing workloads from an SAP HANA system with the Workload Analyzer can help you identify the root cause of performance issues.

### 4.5.2.3.5.1 Analyze Workloads

You can analyze workloads from an SAP HANA system.

#### Prerequisites

- You have the privileges granted by the role `sap.hana.workloadanalyzer.roles::Operator`. You can assign roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Analyze Workload](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it again from the tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

#### Context

The [Analyze Workload](#) app provides you with an overview of the system's health and helps you identify the root cause of performance issues.

#### Procedure

1. Open the [Analyze Workload](#) app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.

The [Analyze Workload](#) app opens displaying two modes, which allow either a real-time or an historical data analysis.

2. The real-time analysis view is displayed by default. If not, click the [Active](#) toggle button to open the real-time analysis.

The real-time analysis view displays two information sets:

- On the upper part of the screen, the chart displays the system resource consumption.

---

Use this chart to identify abnormal conditions. You can customize the information displayed on the load graph in several ways, for example:

- Specify the services
- Specify the refresh period
- Filter the displayed KPIs

For a list of all available KPIs, see *Key Performance Indicators*.

- On the lower part of the screen, the main analysis chart is composed of a main chart and two bar charts. The main chart displays the course of a selected dimension over a given period of time. The bar charts display the top five items within the specified dimensions.

Use the combo boxes to select the needed dimension for each of these charts.

3. Click the *Historical* toggle button to open the historical analysis.

The historical view displays two information sets:

- On the upper part of the screen, the chart displays the system resource consumption. Use this chart to analyze events that happened in the past. You can customize the information displayed on the load graph in several ways, for example:
  - Specify the services
  - Specify the time period  
Use the navigation buttons to move the current time frame by the specified period.
  - Filter the displayed KPIs  
For a list of all available KPIs, see *Key Performance Indicators*.
- On the lower part of the screen the tabular view displays the historical data.

## Related Information

[Deploy a Delivery Unit Archive \(\\*.tgz\) \[page 632\]](#)

[Assign Roles to a User \[page 717\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Key Performance Indicators \[page 314\]](#)

## 4.6 Managing Tables

### Related Information

[Table Types in SAP HANA \[page 332\]](#)

[Basic Table Management in SAP HANA Studio \[page 335\]](#)

[Table and Catalog Consistency Checks \[page 348\]](#)

[Memory Management in the Column Store \[page 352\]](#)

[The Delta Merge Operation \[page 356\]](#)

[Data Compression in the Column Store \[page 368\]](#)

[Table Partitioning \[page 372\]](#)

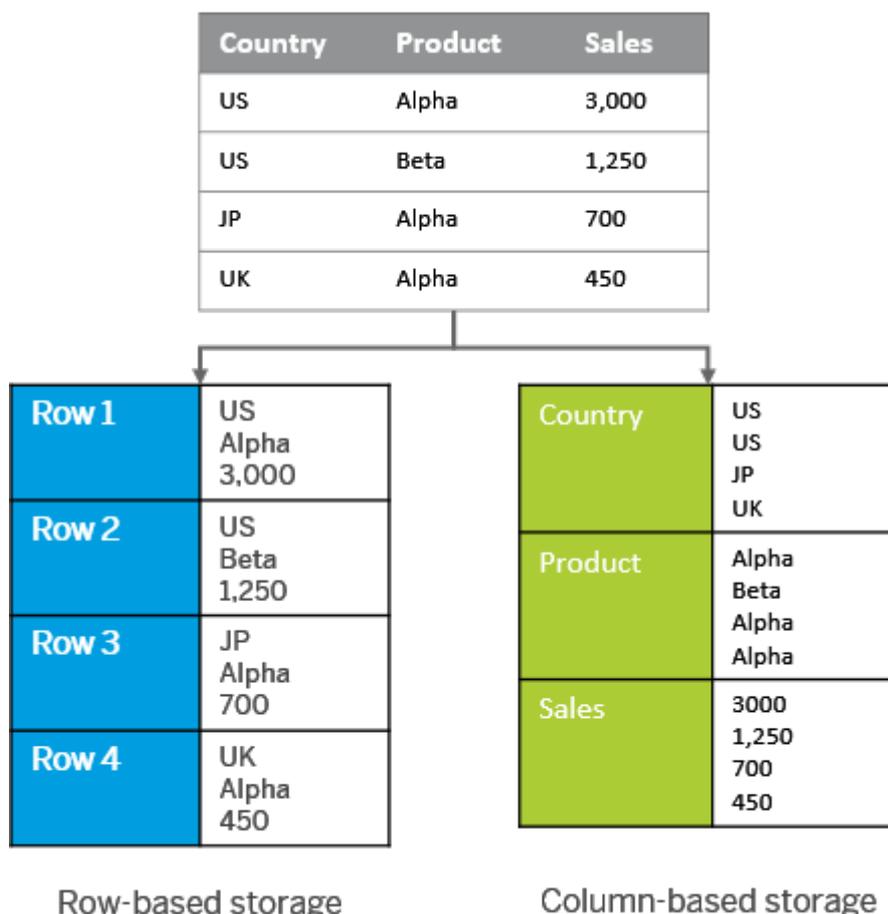
[Table Replication \[page 398\]](#)

[Table Placement \[page 414\]](#)

## 4.6.1 Table Types in SAP HANA

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.

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"> <li>• Calculations are typically executed on individual or a small number of columns.</li> <li>• The table is searched based on the values of a few columns.</li> <li>• The table has a large number of columns.</li> <li>• The table has a large number of rows and columnar operations are required (aggregate, scan, and so on)</li> <li>• High compression rates can be achieved because the majority of the columns contain only a few distinct values (compared to the number of rows).</li> </ul>
Row store	<ul style="list-style-type: none"> <li>• The application needs to process only one single record at one time (many selects and /or updates of single records).</li> <li>• The application typically needs to access the complete record.</li> <li>• The columns contain mainly distinct values so compression rate would be low.</li> <li>• Neither aggregations nor fast searching are required.</li> <li>• The table has a small number of rows (for example, configuration tables).</li> </ul>

#### **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.

### 4.6.1.1 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](#)

---

## 4.6.2 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 \[page 335\]](#)

[Viewing Options for Tables and Views \[page 337\]](#)

[Export Tables and Other Catalog Objects \[page 342\]](#)

[Import Tables and Other Catalog Objects \[page 345\]](#)

[Import ESRI Shapefiles \[page 347\]](#)

[Create a Table in Runtime \[page 340\]](#)

[Create a View in Runtime \[page 341\]](#)

### 4.6.2.1 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.

#### 4.6.2.1.1 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 \[page 337\]](#)

[Table/View Content \[page 339\]](#)

[Data Preview \[page 339\]](#)

### 4.6.2.1.2 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 \[page 337\]](#)

[Table/View Content \[page 339\]](#)

[Data Preview \[page 339\]](#)

## 4.6.2.2 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 \[page 337\]](#)
- [Table/View Content \[page 339\]](#)
- [Data Preview \[page 339\]](#)

### 4.6.2.2.1 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
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 (%)	<p>The current compression ratio of the column in main storage</p> <div style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;"> <p><b>i Note</b></p> <p>If you want to analyze the compression ratio of a table, it must be fully loaded into memory.</p> </div>

## 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 \[page 272\]](#)

### 4.6.2.2.2 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...**

#### Note

LOB values are not formatted. Any LOB data that cannot be visualized is not changed.

### 4.6.2.2.3 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.

## 4.6.2.3 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

#### 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.

## 4.6.2.4 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.

### 4.6.2.5 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
<b>Find objects for export</b>	<ol style="list-style-type: none"> <li>1. From the main menu, choose <b>File &gt; Export</b>.</li> <li>2. Select the export destination <b>SAP HANA &gt; Catalog Objects</b> and choose <i>Next</i>. The <i>Export</i> dialog box appears.</li> <li>3. Search for the objects you want to export and add them to list of objects to be exported on the right.</li> </ol>
<b>Select object for export</b>	<ol style="list-style-type: none"> <li>1. In the <i>Systems</i> view, select the objects that you want to export.</li> <li>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.</li> <li>3. Optional: Search for additional objects to be exported and add them to list of objects to be exported on the right.</li> </ol>

2. Choose *Next*.
3. Specify the scope of the export by choosing the relevant options:

Option	Description
<b>Column Table Format</b>	<p>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.</p> <div style="background-color: #fff9c4; padding: 10px; margin-top: 10px;"> <p><b>i Note</b></p> <p>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.</p> </div>
<b>Export Catalog Objects</b>	<p>Using the following options, you can configure the scope of the export:</p> <ul style="list-style-type: none"> <li>○ 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.</li> <li>○ Select <i>Including Dependencies</i> to export dependent objects of selected objects, that is triggers and indexes. This option is selected by default.</li> </ul>

4. Specify the location to which the file is to be exported:

Option	Description
<b>Export Catalog Objects on Server</b>	<p>The selected catalog objects are saved to a directory on the database server file system. The default directory is <code>/usr/sap/&lt;SID&gt;&lt;instance&gt;/work</code>.</p> <div style="background-color: #fff9c4; padding: 10px; margin-top: 10px;"> <p><b>i Note</b></p> <p>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.</p> </div>

Option	Description
<b>Export Catalog Objects to Current Client</b>	<p>The selected catalog objects are saved in the specified directory on the client file system.</p> <p><b>i Note</b></p> <p>The specified directory must be empty. You can specify a directory that does not exist; it will be created when you start the export.</p>

### ➔ 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](#) ▸).

## 4.6.2.6 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.
- 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
<b>Find objects for import</b>	<ol style="list-style-type: none"><li>1. From the main menu, choose <b>File &gt; Import</b>.</li><li>2. Select the import source <b>SAP HANA &gt; Catalog Objects</b> and choose <i>Next</i>. The <i>Import</i> dialog box appears.</li><li>3. Specify the location of previously exported objects:<ul style="list-style-type: none"><li>○ 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.</li><li>○ 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.</li></ul></li><li>4. Choose <i>Next</i>.</li><li>5. 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).</li></ol>
<b>Select objects for import</b>	<ol style="list-style-type: none"><li>1. In the <i>Systems</i> view, select the objects that you want to import.</li><li>2. From the context menu, choose <i>Import</i>. The <i>Import</i> dialog box appears.</li></ol>

Option	Description
	<p>3. Specify the location of previously exported objects:</p> <ul style="list-style-type: none"> <li>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.</li> <li>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.</li> </ul> <p>4. Choose <i>Next</i>.</p> <p>5. 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).</p>

- Choose *Next*.
- Specify the scope of the import by choosing the relevant options:

Option	Description
<b>Including Data</b>	<p>Object definitions and data are imported.</p> <p><b>i Note</b> This option is selected by default. However, the export must have included both definition and data.</p> <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>
<b>Including Dependencies</b>	Dependent objects of selected objects are also imported, that is triggers and indexes. This option is selected by default.
<b>Replace Existing Catalog Objects</b>	If the objects already exist, they are overwritten.

- 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).

- 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 ▶](#)).

### 4.6.2.7 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* >).

### 4.6.3 Table and Catalog Consistency Checks

Using stored procedures 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, or if possible, you can run the checks on a system copy.

- [Table Consistency Check \[page 349\]](#)
- [Catalog Consistency Check \[page 351\]](#)

## Related Information

[SAP Note 1977584 \(Technical Consistency Checks for SAP HANA Databases\)](#) 

[SAP Note 2116157 - FAQ: SAP HANA Consistency Checks and Corruptions.](#) 

### 4.6.3.1 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.

### 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 - Technical Consistency Checks for SAP HANA Databases](#)

[SAP Note 2116157 - FAQ: SAP HANA Consistency Checks and Corruptions.](#)

## 4.6.3.2 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)
```

## 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.

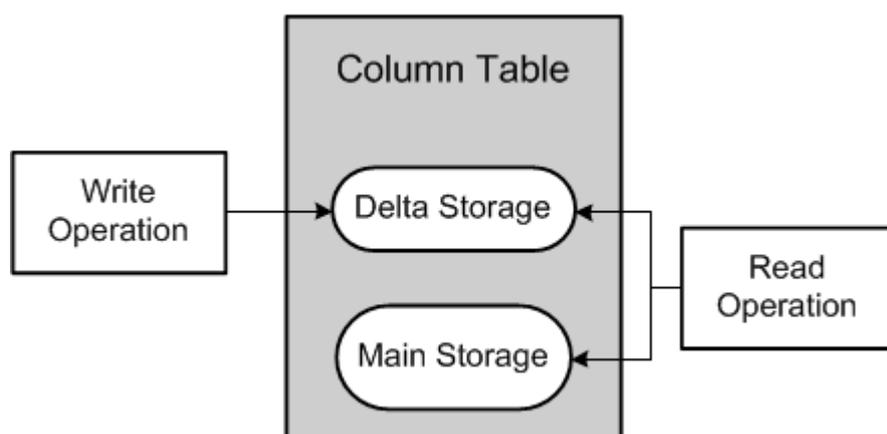
## 4.6.4 Memory Management in the Column Store

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.

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 \[page 356\]](#)

### 4.6.4.1 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 independently, 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, you should normally not have to interfere with this process. However, you can manually load and unload individual tables and table columns if 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)

#### **i** Note

You can see detailed information about a table's current memory usage and load status by viewing its table definition (*Runtime Information* tab).

## Load and Unload a Table from the Systems View

### Procedure

1. In the *Systems* view, navigate to the table.
2. In the context menu of the table, choose *Load into Memory* or *Unload from Memory* as required.
3. Choose *OK*.

---

## 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

Open the SQL console and execute the required statement:

- `LOAD <table_name>`
- `UNLOAD <table_name>`

### 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 an Individual Column Using SQL

### Procedure

Open the SQL console and execute the statement: `LOAD <table_name> (<column_name>, ...)`

### Results

The entire column is loaded or unloaded into or from main memory. Its load status is TRUE or FALSE. The table's load status is PARTIALLY.

## Related Information

[Memory Sizing \[page 272\]](#)

[Memory Management in the Column Store \[page 352\]](#)

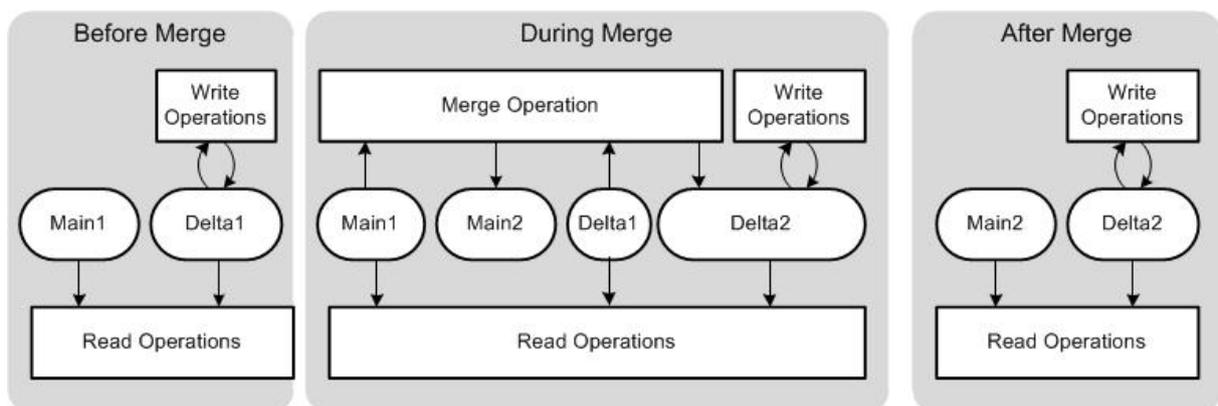
[Table Definition \[page 337\]](#)

### 4.6.5 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.

### 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

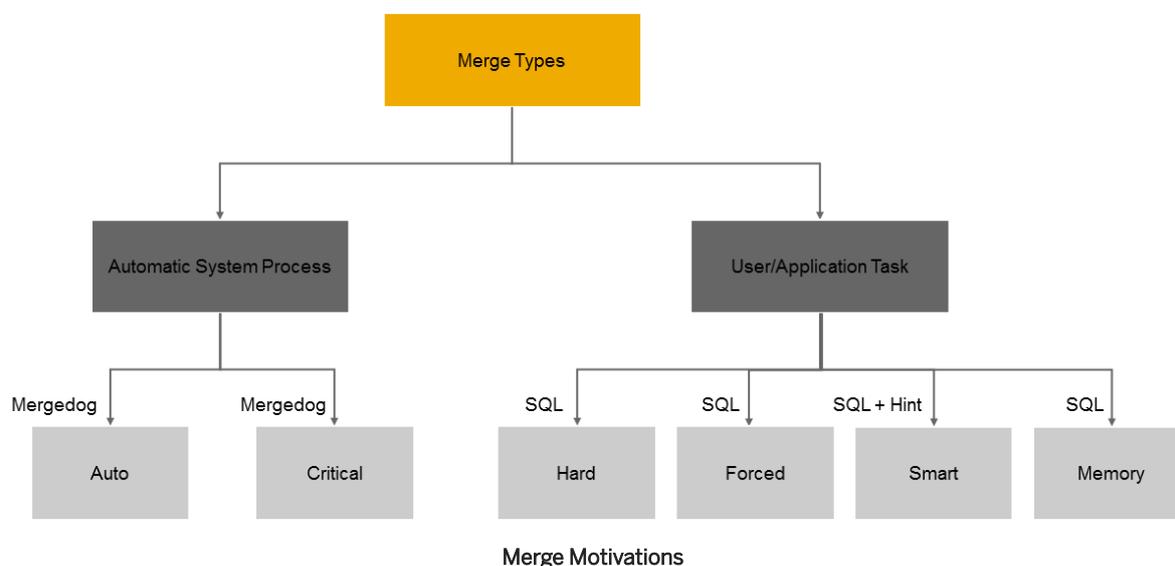
Before a table is (re-)partitioned, a delta merge operation is executed. Therefore, in the case of huge tables, you have to partition them in good time so as not to run out of memory during the merge operation.

## Related Information

## 4.6.5.1 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.

### 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')`.

### 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.

### Note

Main Memory merge was available until SPS 08. See SAP Note - 2057046 for more information.

## 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 \[page 366\]](#)

[Compress a Column Table Manually \[page 371\]](#)

[SAP Note 2057046 - FAQ: SAP HANA Delta Merges](#)

### 4.6.5.2 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.

### 4.6.5.3 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.

Parameter	Meaning
DRC	Delta row count [million] This refers to the current number of rows in the delta storage of the table.
QDW	Queuing delay wait [sec] This refers to the time that a merge thread has been waiting for the merge monitor to allocate it 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&gt;1000 or TMD&gt;3601 or DCC&gt;800 or DMR&gt;0.2*MRC or DLS&gt;5000</pre>	<p>An automatic delta merge of a table is executed if :</p> <ul style="list-style-type: none"> <li>• The size of its delta storage exceeds 1000 MB, or</li> <li>• It has not been merged in over 60 minutes, or</li> <li>• Its delta cell count exceeds 800 million, or</li> <li>• More than 20% of the records in its main storage were deleted, or</li> <li>• The size of its delta log is greater than 5000 MB</li> </ul>

Cost Function Configuration	Meaning
<pre>auto_merge_decision_func = DMS &gt; 1000 or DMS &gt; 42 and weekday(now())=6 and secondtime(now())&gt;secondtime('01:00') and secondtime(now())&lt;secondtime('02:00')</pre>	<p>An automatic delta merge of a table is executed if:</p> <ul style="list-style-type: none"> <li>• The size of its delta storage exceeds 1000 MB, unless</li> <li>• It is Saturday between 1.00 and 2.00, in which case it will be merged if delta storage exceeds 42MB</li> </ul> <p>Note the week starts with Monday as day 0.</p>
<ul style="list-style-type: none"> <li>• <code>smart_merge_decision_func = DMS&gt;1000 or DCC&gt;800 or DLS&gt;5000</code></li> <li>• <code>smart_merge_priority_func = DMS/1000</code></li> </ul>	<p>A delta merge request of a table triggered by a smart merge hint is executed if:</p> <ul style="list-style-type: none"> <li>• The delta storage size exceed 1000 MB, or</li> <li>• The delta cell count in the delta storage is greater than 800 million, or</li> <li>• The size of the delta log is greater than 5000 MB</li> </ul> <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>
<pre>hard_merge_priority_func = QDW</pre>	<p>Delta merges triggered by hard merge are prioritized only by queuing delay wait, in other words, on a first in first out basis.</p>
<pre>hard_merge_priority_function = 1/(7+MMS)</pre>	<p>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.</p>

As of SPS 09 cost functions are also used to optimize compression. For more information see *Data Compression in the Column Store*

## Related Information

[Change a System Property \[page 217\]](#)

[Data Compression in the Column Store \[page 368\]](#)

[SAP Note 2057046](#)

### 4.6.5.4 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

[View Diagnosis Files in SAP HANA Studio \[page 461\]](#)

### 4.6.5.5 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 Administration editor on the [System Information](#) tab.

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"><li>• MERGE for an actual delta merge operation</li><li>• HINT for a merge hint sent to SAP HANA by an application</li><li>• SPARSE for the post-merge optimization of main storage compression</li></ul>
MOTIVATION	This column identifies the underlying merge motivation: AUTO, SMART, HARD, or FORCE

Column	Description
SUCCESS	<p>This column depends on the entry in the TYPE column.</p> <ul style="list-style-type: none"> <li>For MERGE or SPARSE entries, it indicates whether or not the merge or compression optimization operation was successful.</li> <li>For HINT entries, it indicates whether or not the hint from the application to merge was accepted.</li> </ul> <p>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.</p> <div style="background-color: #fff9c4; padding: 10px; margin-top: 10px;"> <p><b>i Note</b></p> <p>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.</p> </div>
LAST_ERROR	This column provides information about error codes of the last errors that occurred (most often 2048). Details are provided in ERROR_DESCRIPTION.
ERROR_DESCRIPTION	<p>The following error codes are possible:</p> <ul style="list-style-type: none"> <li>Error 2480: The table in question is already being merged.</li> <li>Error 2481: There are already other smart merge requests for this table in the queue.</li> <li>Error 2482: The delta storage is empty or the evaluation of the smart merge cost function indicated that a merge is not necessary.</li> <li>Error 2483: Smart merge is not active (parameter <code>smart_merge_enabled=no</code>).</li> <li>Error 2484: Memory required to optimize table exceeds heap limit (for failed compression optimization operations (TYPE=SPARSE, SUCCESS=FALSE)).</li> </ul>
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)

### **i 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.

- 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'
```
- 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 Diagnosis Files in SAP HANA Studio \[page 461\]](#)

[Configure Traces in SAP HANA Studio \[page 465\]](#)

## 4.6.5.6 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
<b>Menu command</b>	<ol style="list-style-type: none"> <li>1. In the <i>Systems</i> view, navigate to the table.</li> <li>2. In the context menu of the table, choose <i>Perform Delta Merge</i>.</li> <li>3. Choose <i>OK</i>.</li> </ol>
<b>SQL</b>	Open the SQL console and execute the required statement: <ul style="list-style-type: none"> <li>◦ <code>MERGE DELTA OF '&lt;table_name&gt;'</code> (hard merge)</li> <li>◦ <code>MERGE DELTA OF '&lt;table_name&gt;' WITH PARAMETERS ('FORCED_MERGE' = 'ON')</code> (forced merge)</li> <li>◦ <code>MERGE DELTA OF '&lt;table_name&gt;' WITH PARAMETERS ('MEMORY_MERGE' = 'ON')</code> (memory-only merge)</li> </ul>

2. Optional: Confirm the delta merge operation in one of the following ways:

- Open the table definition in the table editor 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.

## 4.6.6 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`.

## 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 &gt; 0.010240 and if(OCRC, max(MRC, OCRC) / min(MRC, OCRC) &gt;= 1.75, 1) and (not RP or (RP and TMD &gt; 86400))</code>	Optimize compression runs if <ul style="list-style-type: none"> <li>• The table contains more than 10240 rows AND</li> <li>• (Optimize compression was never run before OR</li> <li>• The number of rows increase or decrease by factor of 1.75)</li> <li>• AND - if range partitioned, the last delta merge happened more than 24 hours ago.</li> </ul>

## Related Information

[SAP Note 1514966](#)

[SAP Note 1637145](#)

[Cost Functions \[page 360\]](#)

## 4.6.6.1 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, in the SAP HANA studio.

### 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 in the SQL console:

```
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:
  - 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 \[page 354\]](#)

### 4.6.6.2 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.

#### Procedure

1. In the Administration editor, open the SQL console.
2. 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.

3. Check the compression status of the table.
4. 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.
5. Check the compression status of the table (see Related Information).

## Related Information

[Check the Compression of a Column Table \[page 370\]](#)

[The Delta Merge Operation \[page 356\]](#)

## 4.6.7 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 DDL statements 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:

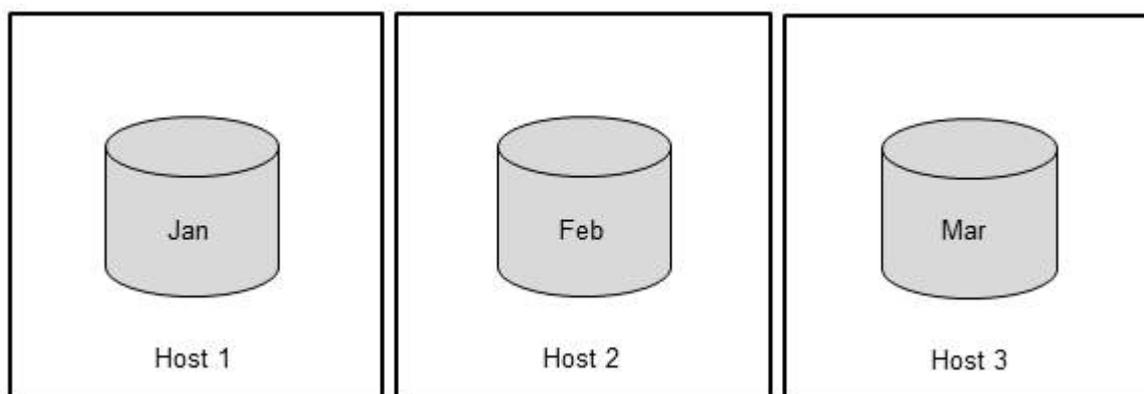
- Load balancing in a distributed system  
Individual partitions can be distributed across multiple hosts. This means that a query on a table is not processed by a single server but by all the servers that host partitions.
- Overcoming the size limitation of column-store tables  
A non-partitioned table 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. If a match is found, it is possible to determine the actual partitions that hold the data being queried. Using this method, the overall load on the system can be reduced, thus improving the response time. For example, if a table is partitioned by year, a query restricted to the data of one year is executed only on the partition with data for this year.

- 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 \[page 356\]](#)

## 4.6.7.1 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 \[page 376\]](#)

### 4.6.7.1.1 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.
<pre>CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT) PARTITION BY ROUNDROBIN 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 or clients that may operate in various landscapes.

## 4.6.7.1.2 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"><li>• 4 partitions on columns a and b are created.</li><li>• The target partition is determined based on the actual values in columns a and b.</li><li>• At least one column has to be specified.</li><li>• If a table has a primary key, all partitioning columns must be part of that key.</li></ul>
<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.

## 4.6.7.1.3 Range Partitioning

Range partitioning creates dedicated partitions for certain values or value ranges in a table. Usually, this requires an in-depth knowledge of the values that are used or valid for the chosen partitioning column. For example, a range partitioning scheme can be chosen to create 1 partition for each calendar month.

Applications may choose to use range partitioning to manage the partitioning of a table actively, that is, partitions may be created or dropped as needed. For example, an application may create a partition for an upcoming month so that new data is inserted into that new partition.

### 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, for example, 1 to 10. 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. If this is not wanted, it is possible to define a rest partition where all rows that do not match any of the defined ranges are inserted. Rest partitions can be created or dropped on-the-fly as desired.

Range partitioning is similar to hash partitioning in that the partitioning column must be part of the primary key. Range partitioning is also restricted in terms of the data types that can be used. See the section Data Types in *Partitioning Limits* for all supported data types.

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

SQL Command	Result
<pre>CREATE COLUMN TABLE MY_TABLE (a INT, b INT, c INT, PRIMARY KEY (a,b)) PARTITION BY RANGE (a) (PARTITION 1 &lt;= VALUES &lt; 5, PARTITION 5 &lt;= VALUES &lt; 20, PARTITION VALUE = 44, PARTITION OTHERS)</pre>	<p>Creates partitions for single values (using =) and for ranges using this semantic: &lt;= VALUES &lt;. This example creates 4 partitions as follows:</p> <ul style="list-style-type: none"> <li>• 1 partition for values greater than or equal to 1 and less than 5</li> <li>• 1 partition for values greater than or equal to 5 and less than 20</li> <li>• 1 partition for values of 44</li> <li>• 1 rest partition for all other values which do not match the specified ranges</li> </ul> <p>The partitioning column has to be part of the primary key. Only STRING, INT and DATE are allowed as data types for the partitioning column.</p>

## Related Information

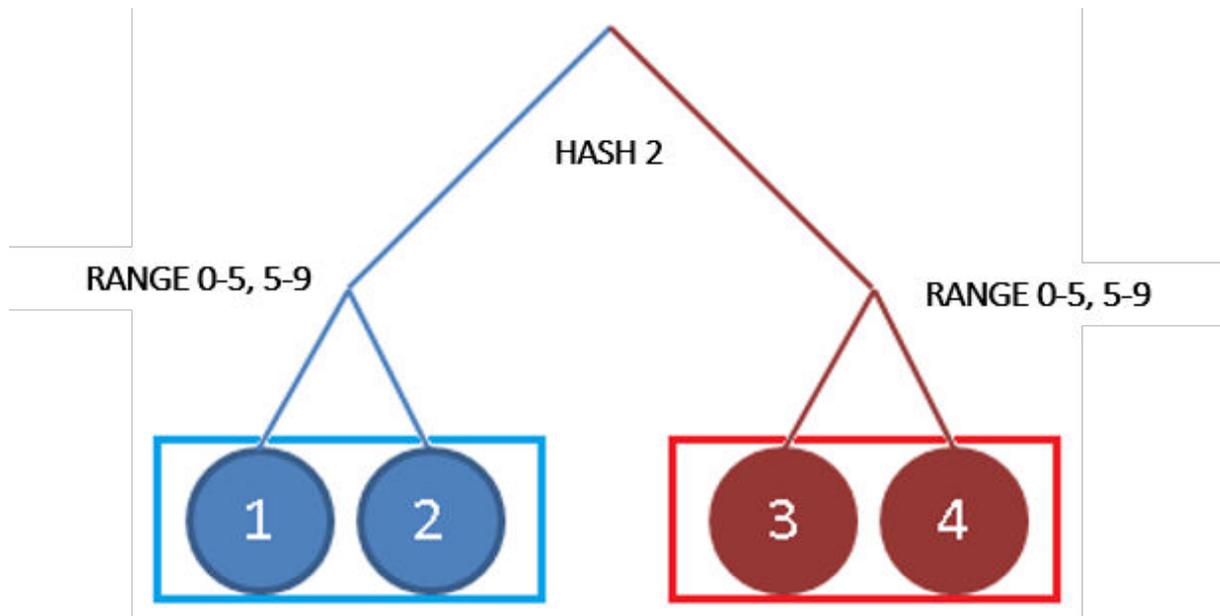
[Partitioning Limits \[page 396\]](#)

### 4.6.7.2 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 illustration shows 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: rows where the value is below 5 and rows where the value is between 5 and 9.

#### Multi-Level Partitioning



The syntax of the SQL code to create these partitions is as follows:

```

PARTITION BY
HASH (<column>) PARTITIONS 2,
RANGE (<column>) (PARTITION 0 <= VALUES < 5, (PARTITION 5 <= VALUES < 9)

```

You can use this approach to implement time-based partitioning, for example, to leverage a date column and build partitions according to month or year.

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 temporal information in its structure, for example a date, multi-level partitioning may be an ideal candidate. If the partitions containing old data are infrequently modified, there is no need for a delta merge on these partitions: the delta merge is only required on new partitions where new data is inserted. Using time-based partitioning in this way the run-time of the delta merge operation remains relatively constant over time as new partitions are being created and used.

As mentioned above, in the second level of partitioning there is a relaxation of the key column restriction (for hash-range, hash-hash and range-range).

When a row is inserted or updated, the unique constraint of the primary key must be checked. If the primary key has to be checked on all partitions across the landscape, this would involve expensive remote calls. Therefore, it is advantageous if only local partitions need to be checked. The concept of partition groups exists for this purpose. It allows inserts to occur whilst only requiring primary key checks on local partitions. All corresponding parts of the second level form a group.

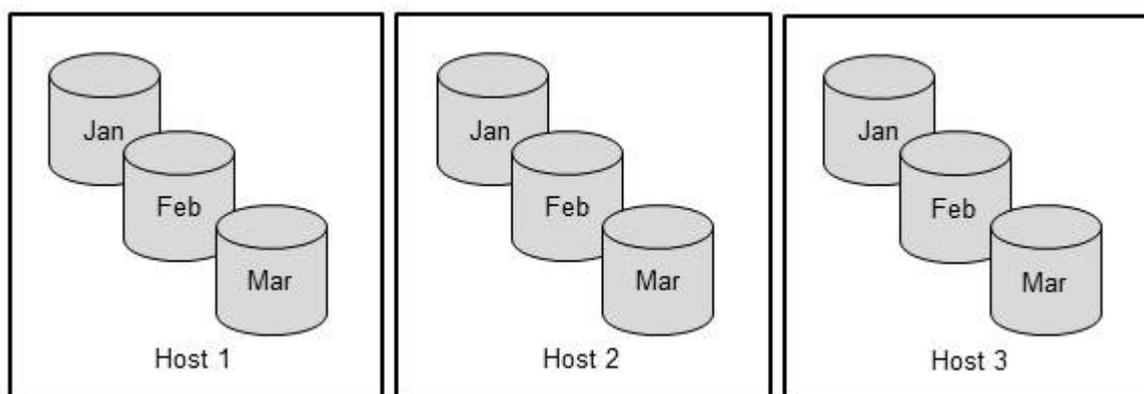
In the figure above, parts 1 and 2 and parts 3 and 4 each form groups. When a row is inserted into part 1, it is only required to check for uniqueness on parts 1 and 2. All parts of a partition group must reside on the same host. When using SQL commands to move partitions, be aware that it is not possible to move individual parts of partition groups, only partition groups as a whole.

## 4.6.7.2.1 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)
```

## 4.6.7.2.2 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)
```

### 4.6.7.2.3 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
```

### 4.6.7.2.4 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)
```

### 4.6.7.3 Using Date Functions to Partition

Use the date functions to partition a table's date or timestamp column by year or month.

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 to restrict your query results by year or by year and month.

## Example

### Creating a Table with Hash Partitioning and Date Function

SQL Command	Result
<pre>CREATE COLUMN TABLE MY_TABLE (a DATE, b INT, PRIMARY KEY (a,b)) PARTITION BY HASH (year(a)) PARTITIONS 4</pre>	<ul style="list-style-type: none"><li>• If a value takes the format "2012-01-08", the hash function is only applied to "2012".</li><li>• This function can also be used for pruning.</li></ul>
<pre>CREATE COLUMN TABLE MY_TABLE (a DATE, b INT, PRIMARY KEY (a,b)) PARTITION BY RANGE (year(a)) (PARTITION '2005' &lt;= values &lt; '2008', PARTITION '2008' &lt;= values &lt; '2011')</pre>	Partition by range using the year.
<pre>CREATE COLUMN TABLE MY_TABLE (a DATE, b INT, PRIMARY KEY (a,b)) PARTITION BY RANGE (month(a)) (PARTITION '2005-01' &lt;= values &lt; '2005-07', PARTITION '2005-07' &lt;= values &lt; '2006-01')</pre>	Partition by range using the year and month.

## 4.6.7.4 Range Partitioning: More Options

Some special features are available for range partitioning: adding additional ranges, deleting ranges, and using dynamic rest partitions.

### 4.6.7.4.1 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.

#### **i** Note

If a partition is created and a rest partition exists, the rows in the rest partition that match the newly-added range are moved to the new partition. If the rest partition is large, this operation may take a long time. If a rest 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 a rest partition. When partitions are dropped, the last partition created cannot be dropped even if a rest 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 rest 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 Rest 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
```

## 4.6.7.4.2 Dynamic Range Partitioning

Dynamic Range Partitioning is available to support the automatic maintenance of the rest partition.

When you create a rest partition there is a risk that over time it could overflow and require further maintenance. Using the dynamic range feature the rest partition will be automatically split into a second partition when it reaches a pre-defined size threshold. You can create partitions with a dynamic rest partition by including the

DYNAMIC keyword in the command when you create the partition. The row count threshold at which automatic splitting takes place is defined either in the SQL syntax or, alternatively, can be set up as an ini file parameter.

### Example

This example creates a rest 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 or a sequence of order numbers.

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;
```

If you define the row count threshold in the SQL command it is saved in the table meta data. As an alternative to this, you can define this value as a system parameter in the `indexserver.ini` file in the `[partitioning]` section. If no value has been specified at table level then this parameter value is used:

```
dynamic_range_default_threshold
```

Enter the value you require, the default value is 10,000,000 rows.

Tables with a dynamic rest partition are monitored by a HANA background job which checks the current size of the partition in comparison to the defined threshold. The background job runs at a predefined interval which is also defined in the `indexserver.ini` file in the `[partitioning]` section in the following initial value:

```
dynamic_range_check_time_interval_sec
```

Enter the number of seconds you require. The default value is 900 seconds. You can deactivate this job by setting the parameter to -1.

## 4.6.7.5 Client-side Statement Routing for Hash Partitioning

If hash partitioning is used on a table, a client tries to connect directly with the server that holds the partitions matching the WHERE clause, if possible.

For example, if a table is partitioned with Hash partitioning on column A and the following SELECT statement is issued then the client will send the request to the node on which the partition is located for which matches to "A = 5" are expected.

```
SELECT * FROM mytable WHERE A = 5
```

This works only if the columns in the WHERE clause match the partitioning columns and are used in expressions with equality ("=").

This reduces the number of hops between the index servers and is especially important if you have a transactional (OLTP) workload. Therefore it should be a design goal to choose a partitioning scheme, that matches your queries. This also works for multi-level partitioning if Hash partitioning is used on the first level. This feature also works for DML.

### **i** Note

For JDBC-based clients, optimal routing happens for prepared statements on Hash partitioned tables. To enable client-side statement routing for Hash partitioning, you must ensure that statement distribution is turned on.

## 4.6.7.6 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 column
  - 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 requires an exclusive lock (only selects are allowed).
- It performs a delta merge in advance.
- It writes everything to the log (required for backup and recovery).

### ➔ 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.

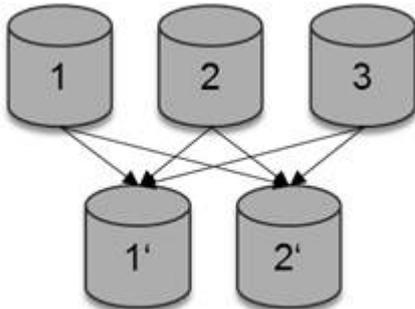
### i Note

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.

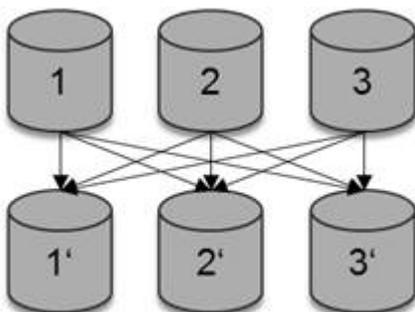
## Re-Partitioning of Partitioned Tables

It is possible to re-partition an already-partitioned table in the following ways:

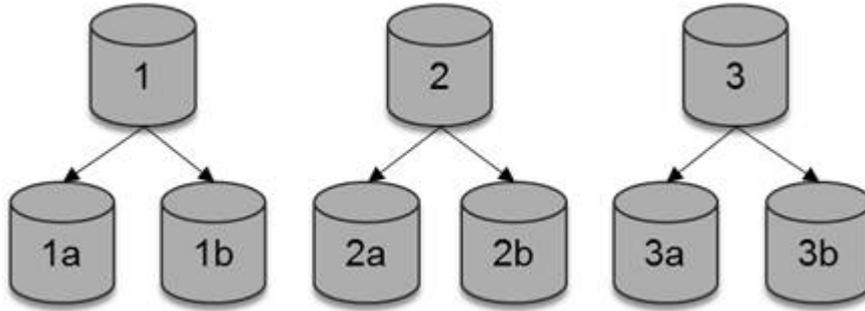
- From  $n$  to  $m$  partitions, where  $m$  is not a multiple/divider of  $n$ , for example, from HASH 3 X to HASH 2 X:



- From  $n$  to  $n$  partitions using a different partition specification or different partitioning columns, for example, HASH 3 X to HASH 3 Y:



- From  $n$  to  $m$  partitions, where  $m$  is a multiple/divider of  $n$ , for example, HASH 3 X to HASH 6 X:



In the first two cases, all source parts must be located on the same host. Up to one thread per column is used to partition or merge the table.

In the third case, it is not necessary to move all parts to the same host. Instead, the partition or merge request is broadcast to each host where partitions reside. Up to one thread per column and source part is used. This type of partitioning operation is typically faster as it is always recommended to choose a multiple or divider of the source parts as the number of target parts. This type of re-partitioning is called a parallel partition/merge.

#### **i** Note

To re-partition a table, use the SQL statement `ALTER TABLE`. For more information, see *SAP HANA SQL and System Views Reference*. You cannot re-partition a table in the SAP HANA studio.

## Parallelism and Memory Consumption

Partitioning operations consume a high amount of memory. To reduce the memory consumption, it is possible to configure the number of threads used. You can change the default value of the parameter `split_threads` in the `partitioning` section of the `indexserver.ini` configuration file. By default, 16 threads are used. In the case of a parallel partition/merge, the individual operations use a total of the configured number of threads for each host. Each operation takes at least one thread.

If a table has a history index, it is possible to partition the main and history index in parallel. Set the parameter `split_history_parallel` in the section `partitioning` in the `indexserver.ini` configuration. The default value is `no`.

## Related Information

[Optimize Table Partitioning \[page 1012\]](#)

[The Delta Merge Operation \[page 356\]](#)

## 4.6.7.6.1 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
<b>Hash</b>	<ol style="list-style-type: none"> <li>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.</li> <li>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.</li> <li>3. Optional: For date or timestamp columns, specify a date function to partition by year or month.</li> </ol>
<b>Round robin</b>	<p>Specify the number of partitions.</p> <p>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.</p>
<b>Range</b>	<ol style="list-style-type: none"> <li>1. 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.</li> <li>2. Optional: For date or timestamp columns, specify a date function to partition by year or month.</li> <li>3. 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. A rest partition is created automatically.</li> <li>4. Choose <i>Validate Input</i> to ensure that the values you have entered are consistent.</li> </ol>

### **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 (|).

6. 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 \[page 1012\]](#)

[Modify Table Distribution Manually \[page 1014\]](#)

### 4.6.7.6.2 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.

### 4.6.7.7 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.

### **i** Note

Tables with time selection partitioning cannot be converted into any other kind of tables using `ALTER TABLE`.

## 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_OBJECTS_DISPOSITION` 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.

## 4.6.7.8 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 in the SQL console of the SAP HANA studio or using the SAP HANA HDBSQL command line tool:

```
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 \[page 349\]](#)

## 4.6.7.9 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 rest 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 a rest partition, a call to ADD PARTITION causes the existing rest partition to be split into a new rest 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 a rest 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, a rest 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 `TablePlacement` 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 a rest 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.

### 4.6.7.9.1 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 a rest partition. Ideally, no rest partition is required.
  - Decide on the number of partitions. Use TablePlacement 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.

### 4.6.7.9.1.1 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.

## 4.6.7.9.2 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.
- Partitioning columns specified for hash and range partitioning must not contain commas (","), dollar signs ("\$\$") or round opening parentheses ("(").
- Ranges must not contain commas (","), minus signs ("-") or asterisks ("\*")
- 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 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. LOB columns are required to be memory LOBs and not disk LOBs.
- Range partitioning: Only the following column store data types are allowed in the partitioning column: TINYINT, SMALLINT, INT, SHORTTEXT, VARCHAR, NVARCHAR, DATE, TIME, TIMESTAMP, SECONDDATE and FIXED. LOB columns must be memory LOBs and not disk LOBs

### **i** Note

Notes on Date Handling and Leap Years:

- The time unit used by partitioning is the whole day, it is not possible to consider smaller units such as hours or minutes.
- Partitioning does not recognize the 29th February as a date. Although the data types DATE, TIME, TIMESTAMP and SECONDDATE are supported (all of these handle leap years correctly) partitioning itself uses CS\_DATE for all date operations which does not handle leap years.

## Partitioning Columns

- The data type of columns used as partitioning columns must not be changed
- The names 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.

---

## Table Replication

- Column Store
  - It is only possible to create replica on all nodes. Table Placement configuration is evaluated though.
  - Tables with history index cannot be replicated
  - Tables have to have a primary key
  - Tables with text columns are not supported
  - Tables for SAP Business Warehouse on SAP HANA must not be replicated.
- Row Store
  - Only row store tables that are located on the master can be replicated via ALTER TABLE.

## Related Information

[SAP Note 2044468 - FAQ: SAP HANA - Partitioning](#)

### 4.6.7.10 Monitoring Partitions

A number of system views allow you to monitor your partitions.

<b>TABLES</b>	contains information on the partition specification.
<b>M_CS_TABLES</b>	Shows run time data per 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.
<b>M_TABLES</b>	Shows row counts and memory usage in an aggregated way for partitioned tables. Information is based on M_CS_TABLES.
<b>M_CS_PARTITIONS</b>	Shows which partitions or sub-partitions form a partition group. This information is for example required when partitions or groups of sub-partitions are to be moved to another host.

## 4.6.8 Table Replication

In a scale-out system tables may be replicated to multiple hosts. This is useful when slowly changing master data often has to be joined with tables or partitions of other tables that are located on multiple hosts and you want to reduce network traffic.

Before using table replication, consider the following aspects:

- Partitioned tables cannot be replicated. For column store tables this means that replicated tables must not exceed the 2 billion row limit.
- Replication of DML statements is executed synchronously, which leads to a performance penalty.

- SAP Business Warehouse (BW) on SAP HANA does not support the use of replicated tables. Do not replicate any tables used in the context of SAP BW on SAP HANA without explicit approval by SAP for the given use cases. Otherwise the use of replicated tables will not be supported
- 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.
- Column store tables have a replica on all available nodes. This is the only supported configuration. For replica creation, the configuration for table placement is taken into consideration.

Table replication must only be used if these considerations are acceptable for your use case.

### Note

- The row store supports replicated tables.
- Only row store tables that are located on the master node can be replicated using ALTER TABLE.
- You cannot replicate column store tables that have history tables, text columns, or column store tables without a primary key.

### Example

#### Creating Column Store Tables with Replicas on All Hosts

SQL Command	Result
CREATE COLUMN TABLE MY_TABLE (I INT PRIMARY KEY) REPLICATA AT ALL LOCATIONS	Creates a column store table with a replica on each available host.
ALTER TABLE MY_TABLE ADD REPLICATA AT ALL LOCATIONS	Replicates an existing column store table on each available host
ALTER TABLE MY_TABLE2 DROP REPLICATA AT ALL LOCATIONS	Drops all replicas

### Example

#### Creating Replicated Row Store Tables

You can use the same syntax as for column store tables to place row store tables on all hosts. Additionally, you can use the following commands to control the location of individual row store replica.

SQL Command	Result
CREATE ROW TABLE MY_TABLE5 (I INT PRIMARY KEY) AT LOCATION ' <u>&lt;master_node&gt;</u> ;	Creates a new row store table with a replica on a specified host (for example, master)
ALTER TABLE MY_TABLE5 ADD REPLICATA AT LOCATION ' <u>&lt;first_slave_node&gt;</u> ';	Replicates an existing row store table at the specified location
ALTER TABLE MY_TABLE5 MOVE REPLICATA FROM ' <u>&lt;first_slave_node&gt;</u> ' TO ' <u>&lt;second_slave_node&gt;</u> ';	Moves an existing replica from one specified location to another

SQL Command	Result
ALTER TABLE MY_TABLE3 DROP REPLICA AT LOCATION '<second_slave_node>;	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

[Table Consistency Check \[page 349\]](#)

[Scaling SAP HANA \[page 998\]](#)

### 4.6.8.1 Asynchronous and Transactional Table Replication

Asynchronous and transactional table replication can help reduce workload on hosts by balancing load across replica tables on worker hosts in a distributed SAP HANA system.

Asynchronous and transactional table 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.

### 4.6.8.1.1 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 \[page 402\]](#)

[Set an Explicit Table Location \[page 403\]](#)

[Asynchronous and Transactional Table Replication Limits \[page 406\]](#)

### 4.6.8.1.1.1 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>
<pre>ALTER SYSTEM ALTER TABLE PLACEMENT (SCHEMA_NAME =&gt; 'REP_SCHEMA') SET (LOCATION =&gt; 'repl_volumes');</pre>	<p>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'.</p> <p>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.</p>

## Example

Create or add a replica table using table placement rules

SQL Command	Result
<pre>CREATE TABLE REP_SCHEMA.TAB1 LIKE SRC_SCHEMA.TAB1 ASYNCHRONOUS REPLICA;</pre>	<p>It creates the first replica table without assigning its location. The replica table will be created on one of replica hosts.</p>
<pre>ALTER TABLE SRC_SCHEMA.TAB1 ADD ASYNCHRONOUS REPLICA;</pre>	<p>This creates more than one replica for SRC_SCHEMA.TAB1</p> <p>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.</p>
<pre>CREATE TABLE REP_SCHEMA.PART_TAB1 LIKE SRC_SCHEMA.PART_TAB1 ASYNCHRONOUS REPLICA;</pre>	<p>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.</p>
<pre>ALTER TABLE SRC_SCHEMA.PART_TAB1 ADD ASYNCHRONOUS REPLICA;</pre>	<p>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.</p>

## 4.6.8.1.1.2 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 REPLICA 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>

SQL Command	Result
<pre>ALTER TABLE SRC_SCHEMA.TAB1 ADD ASYNCHRONOUS REPLICAS AT 'host2:port2'; ALTER TABLE SRC_SCHEMA.TAB1 ADD ASYNCHRONOUS REPLICAS AT 'host3:port3';</pre>	Additional replica tables are created at the specified locations.
<pre>CREATE TABLE REP_SCHEMA.PART_TAB1 LIKE SRC_SCHEMA.PART_TAB1 ASYNCHRONOUS REPLICAS AT ('host1:port1', 'host2:port2', ...);</pre>	For partitioned table, this command creates the first replica table for a partitioned table. The replica partitions will be distributed on the specified nodes.
<pre>ALTER TABLE SRC_SCHEMA.PART_TAB1 ADD ASYNCHRONOUS REPLICAS AT ('host3:port3', 'host4:port4', ...);</pre>	For partitioned table, this command creates additional replica tables on other replica nodes. The replica partitions will be distributed on the specified nodes.

## 4.6.8.1.2 Asynchronous and Transactional 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 in a round robin manner. It is also possible to select a specific replica (identified by volume id) or to use query hints. The following examples illustrate these methods:

#### Round Robin

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 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 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 SQL Hint

You can query replica tables with SQL hint as shown in the following example. By default, a preconfigured hint class exists for Asynchronous Table replication, which is called `hint_result_hana_atr`.

```
SELECT * FROM SRC_SCHEMA.TAB1 WITH HINT(RESULT_LAG('hana_atr'));
```

## Add More Replica Tables to an Active Asynchronous and Transactional 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 REPLICATION;
```

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.

### 4.6.8.1.3 Asynchronous and Transactional Table Replication Limits

General restrictions that apply to the use of Asynchronous and Transactional Table Replication are explained here.

- Asynchronous 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 and Extended Storage table.
- Source tables can be distributed to multiple nodes but a source table and its replica cannot be located at the same node.
- Replica tables cannot be created on the master indexserver.
- Only one identical replica table (or table partition) can exist in a replica node.
- A source table and its replica table should have identical table structure. 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.
- It is not allowed for a “DDL autocommit off” transaction to execute a write operation on a replication source table after a DDL operation on any table in the same transaction boundary.

## 4.6.8.1.4 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'
```

### **i** 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>'.

## 4.6.8.1.4.1 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
```

## 4.6.8.2 Synchronous Table Replication

Synchronous table replication (STR) is a transparent solution without any SQL or application changes. The table replication happens at commit time.

### Differences between Synchronous and Asynchronous Table Replication

With Asynchronous Table Replication (ATR) 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 is a more transparent solution without any SQL or application changes and with most of the benefits of ATR.

Comparing the table replication types

Synchronous Table Replication (STR)	Asynchronous Table Replication (ATR)
<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>

### 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 and Transactional Table Replication Operations* for details.

### Related Information

[Configure Synchronous Table Replication \[page 412\]](#)

## 4.6.8.2.1 Configure Synchronous Table Replication

You can configure synchronous and transactional table replication using the SQL editor by adding replica tables and activating replication.

### Context

In the following steps, SRC\_SCHEMA, SRC\_TABLE and SRC\_PART\_TABLE indicate a normal table's schema name, table name and partitioned table name respectively.

REP\_SCHEMA, REP\_TABLE, and REP\_PART\_TABLE indicate a replica table's schema name, table name and partitioned table name respectively. In addition we assume that SRC\_SCHEMA, REP\_SCHEMA and SRC\_TABLE and SRC\_PART\_TABLE already exist in your system.

### Procedure

1. Add Replica Tables

Normal Table	Partitioned Table
<p><b>Explicit synchronous table creation</b></p> <pre>CREATE COLUMN TABLE REP_SCHEMA.REP_TABLE LIKE SRC_SCHEMA.SRC_TABLE SYNCHRONOUS REPLICA AT 'host:port'</pre>	<p>Explicit synchronous table creation</p> <pre>CREATE COLUMN TABLE REP_SCHEMA.REP_PART_TABLE LIKE SRC_SCHEMA.SRC_PART_TABLE SYNCHRONOUS REPLICA AT 'host:port'</pre>
<p><b>Implicit synchronous table creation</b></p> <pre>ALTER TABLE SRC_SCHEMA.SRC_TABLE ADD SYNCHRONOUS REPLICA AT 'host:port'</pre>	<p>Implicit synchronous table creation</p> <pre>ALTER TABLE SRC_SCHEMA.SRC_PART_TABLE ADD SYNCHRONOUS REPLICA AT 'host:port';</pre>
<p><b>This creates a replica table at the specified location 'host:port'. Note that the replica table cannot be created on a master indexserver and it's not allowed to create multiple replica tables on the same replica node.</b></p>	<p>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.</p>

2. Check Replica Tables

The monitoring view M\_TABLE\_REPLICAS shows general information on synchronous table replicas.

```
SELECT * FROM M_TABLE_REPLICAS WHERE SOURCE_TABLE_NAME = 'SRC_TABLE'
```

- Optional: Deactivate table replication.

Deactivate all synchronous tables:

```
ALTER SYSTEM DISABLE ALL SYNCHRONOUS TABLE REPLICAS
```

Deactivate a specific synchronous table:

```
ALTER TABLE SRC_SCHEMA.SRC_TABLE DISABLE SYNCHRONOUS TABLE REPLICA
```

The above commands deactivate all synchronous replication or a specific synchronous replica's sync. The status of the table's replication is shown in M\_TABLE\_REPLICAS.

- Optional: Activate Replication for all or specific synchronous tables again.

Activate all synchronous tables:

```
ALTER SYSTEM ENABLE ALL SYNCHRONOUS TABLE REPLICAS
```

Activate specific synchronous table:

```
ALTER TABLE SRC_SCHEMA.SRC_TABLE ENABLE SYNCHRONOUS TABLE REPLICA
```

- Optional: Drop replica tables

### **i** Note

If a source table is dropped, its corresponding replica tables are all dropped as well.

Drop all replica tables of the specified source table "SRC\_TABLE":

```
ALTER TABLE SRC_SCHEMA.SRC_TABLE DROP REPLICA AT ALL LOCATIONS
```

Drop 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>
```

You don't have to change anything including queries to access to replica tables. When you access to a source table, one of both source table and replica tables is automatically picked and it is routed to the picked 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.

### **i** Note

The source table should be located at master node. Additionally the same limitations apply to STR as ATR. For more information, see *Asynchronous and Transactional Table Replication Limits*

## Related Information

[Asynchronous and Transactional Table Replication Limits \[page 406\]](#)

## 4.6.9 Table Placement

Using table classification and placement you can control the number of level 1 partitions of tables and the distribution of tables to the hosts of a distributed SAP HANA database.

### Table Classification and Placement

Usually, one table is not the sole storage location for data from a particular application, but is instead closely related to other tables that are also associated with that application.

Table classification enables the table redistribution process to recognize similar or associated database tables. This allows the table redistribution to optimally distribute tables of this kind to different hosts in a distributed SAP HANA database. In this way, you can prevent tables that are regularly joined from being stored on separate hosts, which would mean that the requests first had to be sent across the network when SQL statements are executed. In this case, the tables are placed on the same host and the join can be locally optimized and executed.

For each table, you can specify a group type, subtype, and group name. You can either do this when creating the table or later:

- `CREATE TABLE ... GROUP TYPE <TYPE> GROUP SUBTYPE <SUBTYPE> GROUP NAME <NAME>;`
- `ALTER TABLE ... <TYPE> GROUP SUBTYPE <SUBTYPE> GROUP NAME <NAME>;`
- `ALTER TABLE ... UNSET GROUP`

For more information, see the *SAP HANA SQL and System Views Reference*.

The table classification is stored in the TABLE\_GROUPS table. There is an entry in this table for all tables for which a group type, subtype, or group name was specified. Any tables that do not have an entry in this table have not yet been classified.

For example, the table classification for the tables of a SAP Business Warehouse (BW) Data Store object with the technical name ZFIGL would be:

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

All of the tables in this example have the group type sap.bw.dso and are therefore identified as tables of SAP BW DataStore objects. The subtype is different, depending on the use of each of the individual tables. The group name is the technical name of the DSO in SAP BW. This allows table redistribution to identify that these tables are associated. This means that all three tables are placed on the same host of the distributed SAP HANA database during the table redistribution process.

#### **i** Note

If Table Placement is not configured correctly before data is loaded during migration/installation of SAP BW on SAP HANA, the data distribution may lead to individual hosts being overloaded. The data may be redistributed later on, but this requires a huge amount of time. Therefore this configuration step is crucial.

## Creating New Tables

The table classification and the table placement rules are taken into account on two occasions: table redistribution and during table creation. In the example above, this would mean that an additional table with the group name ZFIGL would be created on the same host as the existing tables.

## Number of Level 1 Partitions

Table classification can also be used to control the number of partitions at level 1 of the partitioning specification. The DSO ZFIGL from the example above has the following partitioning:

- Level 1: HASH (DOCNR, LINEITEM)
- Level 2: RANGE (CALMONTH)

The number of partitions at level 2 (range for the CALMONTH column) is defined and managed by the application. The number of partitions at level 1 and therefore the distribution of a table across multiple hosts, on the other hand, can be determined by the table redistribution. For example, small tables should only have one partition here, and should not be distributed. Very large tables, on the other hand, should be stored with as many partitions as possible on a large number of hosts.

## Customizing Table Placement

The rules that are used for table placement and for determining the level 1 partitions are stored in the view TABLE\_PLACEMENT in the schema \_SYS.

The TABLE\_PLACEMENT view has the following columns for specifying groups of tables or individual tables: SCHEMA\_NAME, TABLE\_NAME, GROUP\_NAME, GROUP\_TYPE, and SUBTYPE

You do not need to fill all of these columns when creating and maintaining rules. A more specific entry takes precedence over a more general entry. For example, if you only fill SCHEMA\_NAME, the rules apply for all tables in the schema. If there is also an entry in which SCHEMA\_NAME and GROUP\_TYPE are filled, this entry overwrites the first, more general entry for the corresponding tables.

For a more detailed overview of the priorities when evaluating the entries in the TABLE\_PLACEMENT view, see SAP Note 1908082.

You can control the number of level 1 partitions with the following columns:

Column	Description
MIN_ROWS_FOR_PARTITIONING	This defines the minimum number of records that must exist in a table before level 1 partitioning takes place.
INITIAL_PARTITIONS	If the threshold value in the MIN_ROWS_FOR_PARTITIONING column is exceeded, the table redistribution performs a partitioning. The initial number of partitions is stored in this column.

Column	Description
REPARTITIONING_THRESHOLD	Once a table has been partitioned with the specified initial number of partitions, for performance reasons, the table is only repartitioned by doubling the number of partitions. For example, if the initial number of partitions is three, this would result in six partitions being created during a repartitioning. You maintain the threshold value for the number of records in a partition that triggers a repartitioning of this kind in this parameter.
LOCATION	The table placement is controlled using the column LOCATION. Possible values are MASTER, SLAVE, and ALL. You can use these to determine the types of nodes of a distributed SAP HANA database on which the respective tables can be stored.

However, as a general principle, the system never creates more partitions than the number of available hosts. For example, if a distributed SAP HANA database only has five hosts, the repartitioning described above, from three partitions to six partitions, would not take place. The partitioning rules also apply only to tables for which Hash or RoundRobin partitioning is specified at level 1. Tables without a partitioning specification are also not automatically partitioned when the threshold values are exceeded.

## Example: Customizing SAP BW InfoCubes

SCHEMA_NAME	GROUP_TYPE	MIN_ROWS_FOR_PARTITIONING	INITIAL_PARTITIONS	REPARTITIONING_THRESHOLD	LOCATION
		2,000,000,000	3	2,000,000,000	slave
SAPBWP		2,000,000,000	3	2,000,000,000	master
SAPBWP	sap.bw.cube	40,000,000	3	40,000,000	slave

### **i** Note

Only relevant columns from the TABLE\_PLACEMENT table are shown above. Empty columns (that is, TABLE\_NAME, SUBTYPE, GROUP\_NAME) were left out for clarity.

The first entry in the TABLE\_PLACEMENT view is the default entry. All specification columns are empty for this entry. This ensures that there is always an applicable rule for creating or distributing tables. The second entry overwrites the location of the default entry (slave) for all tables of the SAP BW schema. The third entry defines the size of level 1 partitions and their location for InfoCubes in the SAP BW schema.

This customizing means that all InfoCube tables are stored on the slave nodes of the distributed SAP HANA system. InfoCube tables with more than 40 million records are divided into three partitions. If the individual partitions have more than 40 million entries on average, these are halved again in accordance with the REPARTITIONING\_THRESHOLD threshold value. This results in a table with six partitions. (As described above, the system does not create more partitions for a table than there are hosts available. The dividing of the three initial partitions into six therefore only takes place, if the SAP HANA database has at least six slave nodes. If there are five or fewer nodes, the number of partitions in this customizing example is limited to three. For

performance reasons, there is no automatic repartitioning from three to five partitions, for example. If you want to repartition to five partitions, you need to enter a corresponding customizing rule in the TABLE\_PLACEMENT view.

There is also no automatic repartitioning when threshold values are exceeded. Instead, this is proposed during the next execution of the table redistribution process.

Tables in the SAP BW schema that do not belong to the group type sap.bw.cube are stored on the master node, in accordance with the second entry. These include, for example, the tables of the SAP Basis component, and those of the ABAP runtime environment.

Tables that are not in the SAP BW schemas (for example, tables replicated using SLT) are always stored on a slave node, in accordance with the first entry.

### **i** Note

For reasons of clarity, this example only takes account of the InfoCube tables of an SAP BW system. For a complete description of the table placement for SAP BW, see SAP Note 1908073.

## Table Redistribution Parameters

The following parameters influence the behavior of the table redistribution in relation to the table classification and placement functionality.

Parameter	Description
Global Config (global.ini): [table_placement] same_num_partitions	If this parameter is set to true (default: false), all tables with the same group name have the same number of level 1 partitions. The number of partitions is determined by the largest table within the group.
Service Config (typically 'indexserver.ini'): [table_redist] all_moves_physical	By default, when the table redistribution is performed, tables are only moved to the new node with their working memory part. The persistence part is written to the new node during the next delta merge. If you want to have the persistence part moved immediately during the table redistribution, set this parameter to true. Note, however, that this can significantly extend the runtime.
[table_redist] force_partnum_to_splitrule	Setting this parameter to true forces the execution of operations that change the number of level 1 partitions. For example, if a table has two level 1 partitions but should have three according to the Customizing settings, the table redistribution process would not, by default, adjust this. Activating this option forces the adjustment.
global.ini: [table_placement] method 2	The classification of row store tables is currently not taken into account. To ensure that these tables are always created on the master node, or are moved there, you need to set the parameter 'method' in the section [table_placement] of global.ini to "2".

---

## Important Notes

- All tables in the 'SYS%' and '\_SYS%' schemas are explicitly excluded from this functionality.
- The customizing of the TABLE\_PLACEMENT table replaces the INI parameters used before SAP HANA SPS 06 for the column store.
- The table redistribution process is performed in two steps:
  1. Generation of a plan
  2. Execution of the plan

## Related Information

[Scaling SAP HANA \[page 998\]](#)

[Table Distribution in SAP HANA \[page 1004\]](#)

[SAP Note 1908082](#)

[SAP Note 1908073](#)

## 4.7 Workload Management

Workload management allows you to determine how much concurrent work takes place on an SAP HANA system and how that work is prioritized.

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.

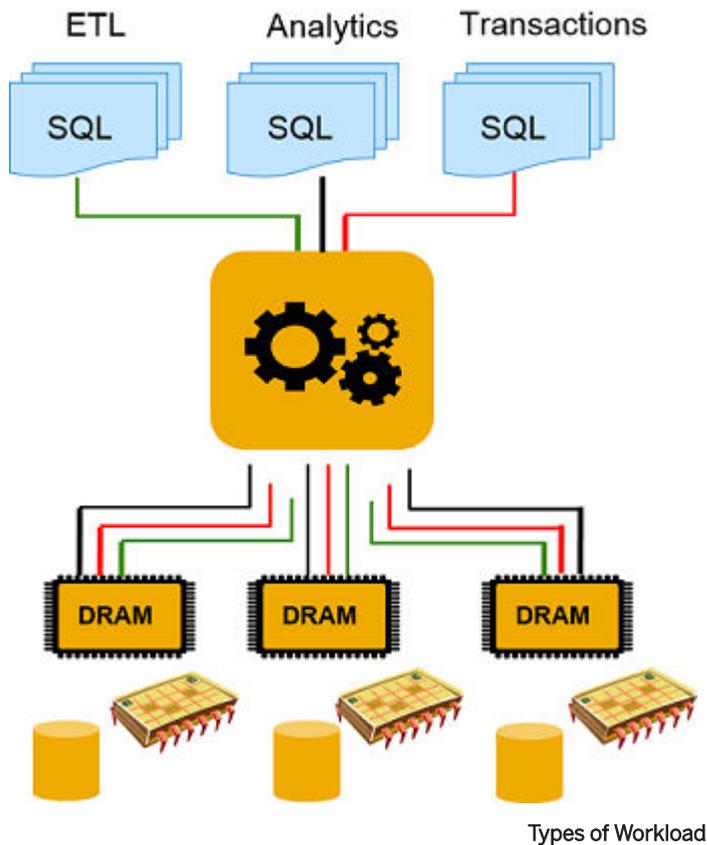
### 4.7.1 Workload in the Context of SAP HANA

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 \[page 264\]](#)

[Scaling SAP HANA \[page 998\]](#)

### 4.7.1.1 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.
- 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
<b>CPU</b> Configure CPU at Operating System level	This approach applies primarily to the use cases of SAP HANA multitenant database containers and multiple SAP HANA instances on one server.  In this case settings related to affinity (binding processes to CPU cores) are applied in <code>daemon.ini</code> . Processes must be restarted before the changes become effective.  For more information, see <i>Controlling CPU Consumption</i> .
<b>CPU</b> 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> .
<b>Memory</b> Change global settings related to 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> .

Area	Possible Actions
<p><b>Priority and Dynamic Workload Class Mapping</b></p> <p>Manage workload and workload priority using classes.</p>	<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 by either:</p> <ul style="list-style-type: none"> <li>• SQL commands</li> <li>• Using the features of the HANA Cockpit.</li> </ul> <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 \[page 424\]](#)

[Controlling Parallel Execution of SQL Statements \[page 427\]](#)

[Setting a Memory Limit for SQL Statements \[page 277\]](#)

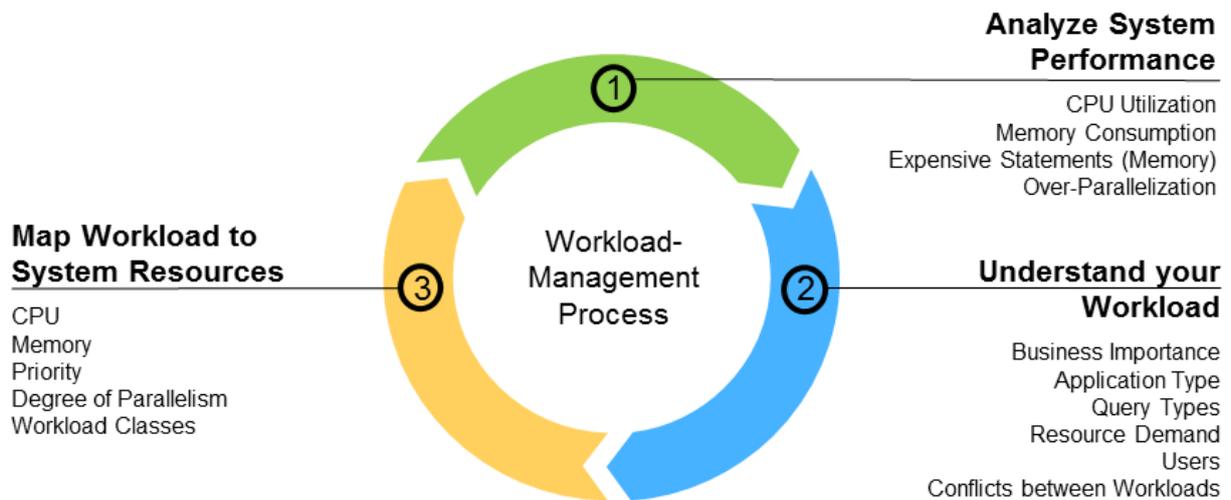
[Managing Workload with Workload Classes \[page 435\]](#)

[Example Workload Management Scenarios \[page 444\]](#)

### 4.7.1.2 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.

### 4.7.1.3 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 \[page 424\]](#)

[Parameter Reference: CPU \[page 429\]](#)

[Controlling Parallel Execution of SQL Statements \[page 427\]](#)

[Parameter Reference: Memory Consumption \[page 281\]](#)

[Managing Workload with Workload Classes \[page 435\]](#)

[Example Workload Management Scenarios \[page 444\]](#)

[SAP Note 2215929](#)

## 4.7.2 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 multitenant database containers 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 `indexserver.ini` configuration file. This may be more convenient and does not require the system to be offline. For more information, see *Managing Resources in Multiple-Container Systems*.

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.

### 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

#	Feature	Example Value
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.

### 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 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 takes effect after a restart of the affected SAP HANA processes.

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 in a MDC) and `numactl` controls NUMA policy for processes or shared memory.

#### 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 setting:

```
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 setting:

```
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 multitenant database container 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';
```

## Related Information

[Managing Resources in Multiple-Container Systems \[page 143\]](#)

## 4.7.3 Controlling Parallel Execution of SQL Statements

You can apply ini file settings to control the two thread pools SQLExecutors and JobExecutors 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:

- **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.

For both SqlExecutor and JobExecutor, a separate limit can be set for the maximum number of threads. For example, on a system where OLAP workload would normally consume too many CPU resources you can apply a maximum value 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 the following subsection, *Parameter Reference: CPU*.

## Parameters for SQL Executor

The following SqlExecutor parameters are in the `sql` section of the `indexserver.ini`.

- `sql_executors` - The target number of threads that can be used.
- `max_sql_executors` - The maximum number of threads that can be used. Not set by default so that new threads are created to handle incoming requests.

## Parameters for JobExecutor

The following JobExecutor parameters are in the `execution` section of the `global.ini` or `indexserver.ini`.

- `max_concurrency` - The target number of threads that can be used.
- `max_concurrency_hint` - Limit concurrency hint even if more active job workers would be available.
- `default_statement_concurrency_limit` - Used to restrict the actual degree of parallel execution per connection within a statement.

JobExecutor settings do not solely affect OLAP workload, but also other SAP HANA subsystems (for example, memory garbage collection, savepoint writes). The JobExecutor executes also operations like table updates and backups, which were delegated by the SqlExecutor. JobExecutor settings are soft limits, meaning the JobExecutor can “loan” threads (this applies to the SQL Executor as well), if available, and then fall back to the maximum number of threads when done.

### Tip

In a system that supports multitenant database containers, a reasonable default value for the `max_concurrency` parameter is the number of cores divided by the number of tenant databases. Do not specify a value of `0`.

The following parameter is in the `parallel` section of the `indexserver.ini`. For details see, *Parameters that Control CPU*.

- `num_cores` - The number of threads (logical cores) available for execution.

## Related Information

[Understand your Workload \[page 421\]](#)

[Example Workload Management Scenarios \[page 444\]](#)

[Parameter Reference: CPU \[page 429\]](#)

### 4.7.3.1 Parameter Reference: CPU

The `sql`, `execution` and `parallel` sections of the `global.ini` or `indexserver.ini` file contain parameters that allow you to control the CPU consumption of SAP HANA.

#### Parameters for SQL Executor

You can change the default settings in the configuration editor of the SAP HANA studio (recommended) or directly in the `sql` section of the `global.ini` or `indexserver.ini` system properties file.

`sql_executors` - sets the target number of threads that can be used.

- The parameter `sql_executors` defines the number of requests that can be processed by the system. This is the target number of threads in the `SqlExecutors` thread pool, and it means the 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. As each thread allocates a particular amount of main memory for the stack, reducing the value of this parameter can help to reduce memory footprint.
- The default value is the number of logical cores in a system (0).
- Does not require a restart. Available since SPS 08.

`max_sql_executors` - sets the maximum number of threads that can be used.

- The parameter `max_sql_executors` defines the maximum number of threads.
- Does not require a restart. Available since SPS 10.

#### Caution

SAP HANA will not accept new incoming requests if the limit is exceeded. As a result handle with extreme care.

#### Parameters for Job Executor

You can change the default settings in the configuration editor of the SAP HANA studio (recommended) or directly in the `execution` section of the `global.ini` or `indexserver.ini` system properties file.

`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.
- Does not require a restart. Available since SPS 09.

`max_concurrency_hint` - limit concurrency hint even if more active job workers would be available.

- The JobExecutor proposes the number of jobs to create for parallel processing based on the recent load on the system. This parameter limits decisions to parallelize on a low level of code. 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 given, but hint is never greater than `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. Performance tests with the target workload are needed to confirm this.
- Does not require a restart. Available since SPS 08 (revision 85).

`default_statement_concurrency_limit` - restrict the actual degree of parallel execution per process within a statement

- If this parameter is used to tune throughput, and `max_concurrency_hint` and `num_cores` are also set explicitly, then we recommend that `num_cores` and `max_concurrency_hint` should be smaller or equal to the value for `default_statement_concurrency_limit`.
- Set this to a reasonable value between 1 and `max_concurrency`. If set to 0 there is no limitation.
- Does not require a restart. Available since SPS 10 (revision 102.3).

### **i** Note

You may need to add this parameter before setting it, if it is not visible in the [Configuration](#) tab of the Administration editor.

`num_cores` - this parameter is in the **parallel** section of the `global.ini` or `indexserver.ini` system properties file.

- Before SPS 09 this parameter was the main configuration to influence the low-level strategy to parallelize execution in SAP HANA. It's advised to set `num_cores` and `max_concurrency_hint` to the same numerical value, because the two values are used in the same logical function, but affect different places in the code.
- Default: Actual setting for `max_concurrency`, uses all available logical cores if nothing else is specified).
- Does not require a restart. Available since SPS 08.

---

## 4.7.4 Setting a Memory Limit for SQL Statements

The statement memory limit allows you to set a limit both per statement and per SAP HANA host.

### Prerequisites

To apply these settings you must have the system privilege INIFILE ADMIN.

### 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.

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.

This limit only applies to single statements, not the system as a whole. Tables which require much more memory than the limit applied here may be loaded into memory.

You can view the (peak) memory consumption of a statement in `M_EXPENSIVE_STATEMENTS.MEMORY_SIZE`.

Note that `M_EXPENSIVE_STATEMENTS.REUSED_MEMORY_SIZE` is not used as of SPS 09.

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.

### 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. In the `global.ini` file, expand the `memorymanager` section and set the parameter `statement_memory_limit`. Set a statement memory limit in GB (integer values only) with a value between 1 and some fraction of the global allocation limit.

### **i** Note

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).

Statements that exceed the limit you have set on a host are stopped by running out of memory.

3. In the `global.ini` file, expand the memorymanager section and set the parameter `statement_memory_limit_threshold` as a percentage of the global allocation limit (`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 the overall `global_allocation_limit` parameter setting.

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.

4. To set a user-specific statement limit and exclude a user from the global limit use the ALTER USER statement as shown here:

```
ALTER USER <user_name> SET PARAMETER STATEMENT MEMORY LIMIT = <gb>
```

- 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)

### **i** Note

Setting the statement memory limit to 0 will disable any statement memory limit for the user, or, to reset a user-specific limit use the CLEAR option:

```
ALTER USER <user_name> 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
<code>global_allocation_limit</code>	Default: 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
<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

[Parameter Reference: Memory Consumption \[page 281\]](#)

### 4.7.4.1 Parameter Reference: Memory Consumption

The `memorymanager` section of the `global.ini` file contains parameters that allow you to control the memory consumption of SAP HANA.

You can change the default settings in the configuration editor of the SAP HANA studio (recommended) or directly in the **global.ini** system properties file.

These parameters require tracking to be enabled in `global.ini` [`resource_tracking`]. Resource tracking creates a runtime overhead in SAP HANA:

- `enable_tracking = on`
- `memory_tracking = on`

### **i** Note

In a system that supports multitenant database containers, you can configure the `global.ini` at both the system level and the database level. Parameters configured at the system level apply to the complete system and all databases. Parameters configured at the database level apply to the specified database only.

`global_allocation_limit` - limits the amount of memory that can be used by the system as a whole.

- The parameter `global_allocation_limit` defines the maximum memory allocation limit in MB.
- The global allocation limit is calculated by default 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. A missing entry or a value of 0 results in the system using the default settings.
- Does not require a restart. Available since SPS 08.

### **i** Note

In a system that supports multitenant database containers, the global allocation limit configured at the system layer of the `global.ini` file is always effective regardless of any value configured at the database layer.

`statement_memory_limit` - defines the maximum memory allocation per statement in GB.

- 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.
- The default value is 0 (no limit). Set this parameter to a value between 1 GB and the value of the global allocation limit.
- Does not require a restart (applies to new statements). Available since SPS 09.

`statement_memory_limit_threshold` - defines the maximum memory allocation per statement as a percentage of the global allocation limit.

- If a value for this parameter has been set then the statement memory limit is only applied if the current SAP HANA memory consumption exceeds the statement memory limit threshold as a percentage of the global allocation limit.
- The default value is 0% (the `statement_memory_limit` is always respected). Set this parameter to a value between 1 GB and the value of the global allocation limit.
- Does not require a restart (applies to new statements). Available since SPS 09.

## Related Information

[Allocated Memory Pools and Allocation Limits \[page 273\]](#)

## 4.7.5 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 (for example, a statement memory limit). Workload classes allow SAP HANA to influence dynamic resource consumption on the session or statement level. A client submits client attributes in the connect string including the 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 following parameters can be set:

- PRIORITY
- STATEMENT MEMORY LIMIT
- STATEMENT THREAD LIMIT

Managing workload classes requires the 'WORKLOAD ADMIN' privilege. Changes of workload classes or mappings will only be applied when a (connected) database client reconnects.

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.

### 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:

Parameter	Value
<code>statement_priority</code>	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.
<code>statement_thread_limit</code>	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.
<code>statement_memory_limit</code>	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.

## 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'
```

## 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:

1. Application User Name: this is usually the user logged into the application.
2. Client: the client number is usually applied by SAP ABAP applications, like SAP Business Suite / BW.
3. Application Component Name: this value is used to identify sub-components of an application, such as CRM inside the SAP Business Suite.
4. 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.
5. Application Name
6. 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.

## Points to Note

- Workload classes also work in a multiple container system, but you need to define workload classes for each tenant database
- Workload classes are applied to the complete SAP HANA database in a scale-out environment, but not to each single node

## Related Information

[Create a Workload Class Mapping \[page 442\]](#)

### 4.7.5.1 Working with Workload Classes: Monitor, Export, Disable

You can use system views to see 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;
```

## Import and Export Class Details

Note that 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}
```

### 4.7.5.2 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.

The following configuration options are possible:

- Create workload classes
- Create workload class mappings
- Create user-specific parameters

In addition, you can perform the following actions on alert checkers:

- Edit global limits
- Enable memory tracking

#### **i** Note

You can configure alerts with the SAP HANA cockpit only if the monitoring and alerting functions in the system are being implemented by the **embedded statistics service**, not the statistics server. For more information about migrating to the statistics service, see SAP Note 1917938.

## Related Information

[Creating a Workload Class \[page 441\]](#)

[Create a Workload Class Mapping \[page 442\]](#)

[Create User-Specific Parameters \[page 443\]](#)

[SAP Note 1917938](#)

### 4.7.5.2.1 Applying 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.

#### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Manage Workload Classes](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it again from the tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

#### Context

The [Workload Classes](#) app 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. For more information, see [Managing Workload](#) in Related Information.

#### Procedure

1. Open the [Workload Classes](#) app by clicking the [Manage Workload Classes](#) tile on the homepage of the SAP HANA cockpit.

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 in the [Monitor Statements](#) app.

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 \[page 435\]](#)

[Tile Catalog: SAP HANA Database Administration \[page 35\]](#)

### 4.7.5.2.2 Creating a Workload Class

You can create workload classes to manage the workload of the SAP HANA system.

## Context

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 the [Workload Classes](#) app by clicking the [Manage Workload Classes](#) tile on the homepage of the SAP HANA cockpit.
2. Select [Create](#). Specify the workload class properties, then select [Create](#).

Field Name	Description
Name	Name of the workload class
Priority	Priority
Statement Memory Limit	Maximum amount of memory used to execute the statement

Field Name	Description
Statement Thread Limit	Maximum number of threads executed by the statement

3. You can also immediately create a mapping for the workload class by entering the mapping properties under *Mapping Details (Optional)*. Refer to the following topic *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.

### 4.7.5.2.3 Create a Workload Class Mapping

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 the *Workload Classes* app by clicking the *Manage Workload Classes* tile on the homepage of the SAP HANA cockpit.
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. 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 ▶ client ▶ application component name ▶ application component type ▶ application name ▶ user name ▶**. For example, a mapping where the application user is matched would take precedence over a mapping where the database user matched.

Field Name	Description
Name	Name of the mapping (grouped by Application User Name)
APPLICATION USER NAME	Name of the application user
CLIENT	Client
APPLICATION COMPONENT NAME	Name of the application component
APPLICATION COMPONENT TYPE	Name of the component type
APPLICATION NAME	Name of the application

Field Name	Description
USER NAME	Name of the database user

## Results

The workload class mapping is created and displayed in the list.

### 4.7.5.2.4 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 the [Workload Classes](#) app by clicking the [Manage Workload Classes](#) tile on the homepage of the SAP HANA cockpit.
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.

## 4.7.6 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
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.

	Analytics Workload	Mixed Workload
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.
<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 \[page 421\]](#)

[Analyzing System Performance \[page 422\]](#)

---

## 4.8 Scheduling of Recurring Administration Tasks

SAP HANA Extended Services (SAP HANA XS) 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 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 \[page 1110\]](#)

## 4.9 Hardware Checks 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 a 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 in *SAP Note 2133369*.

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

[Product Availability Matrix for SAP HANA](#) 

[SAP Note 1943937 - Hardware Configuration Check Tool - Central Note.](#) 

[SAP Note 2133369](#) 

## 4.9.1 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 - Hardware Configuration Check Tool - Central Note*.
- Check *SAP Note 2235581* for supported OS 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 should work in both directions for fully qualified domain names.
- SAPCAR is needed to extract the binaries.

---

## Context

Follow the instructions in *SAP Note 1943937 - Hardware Configuration Check Tool - Central Note* to download the latest version of the tool as a SAR file from the SAP Service Marketplace.

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 Linux 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 - Hardware Configuration Check Tool - Central Note*.

## Related Information

[SAP Note 1943937 - Hardware Configuration Check Tool - Central Note](#)

[SAP Note 2235581](#)

[SAP Note 2298750](#)

[SAP Software Download Center](#)

## 4.10 SAP Solution Manager for SAP HANA

You can integrate SAP HANA into an overall operations concept supported through SAP Solution Manager

SAP HANA and SAP HANA-based applications are supported by SAP Solution Manager. If you already use SAP Solution Manager, the effort for integrating SAP HANA into your existing operations concept is relatively low. You can use its functions for SAP HANA related activities.

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.

---

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

## Related Information

[SAP Note 1747682: Managed System Setup for HANA](#)

### 4.10.1 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, SP05.

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

[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](#)

---

## 4.10.2 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 afterwards.

### Related Information

[Configure SLD Registration Using the Graphical User Interface \[page 451\]](#)

[Configure SLD Registration Using the Command-Line Interface \[page 453\]](#)

[Configure SLD Registration Using the Web User Interface \[page 455\]](#)

### 4.10.2.1 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>/hdblcmm
```

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

2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

```
./hdblcmmgui
```

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.
<i>Use HTTPS</i>	Specifies whether or not to use HTTPS.

5. 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
```

## 4.10.2.2 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:

- `./hdbclm`

and enter the index of the `configure_sld` action, or

- 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.

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</i>	Specifies whether or not to use HTTPS.

For more information about parameters for SLD configuration, see Related Information.

- 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* in Related Information.

### 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=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
```

## 4.10.2.3 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](#).
    - 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* in Related Information.

- You are logged on as the system administrator user `<sid>adm`.

You should verify that the following additional prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from the SAP HANA studio.

- The SAP HANA studio revision is 120 or higher.
- For Linux:
  - The system property `org.eclipse.swt.browser.XULRunnerPath` should be set in `hdbstudio.ini` to point to the path of XULRunner, for example:  
`-Dorg.eclipse.swt.browser.XULRunnerPath=<path to xulrunner>`.  
This `hdbstudio.ini` file is located in the same folder as the executable that is used to start the SAP HANA studio. For Linux, the default location is `hana/shared/<SID>/hdbstudio..`

## 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
<b>Web Browser</b>	<p>Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser:</p> <p><code>https://&lt;hostname&gt;:1129/lmsl/HDBLCM/&lt;SID&gt;/index.html</code></p> <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>The URL is case sensitive. Make sure you enter upper and lower case letters correctly.</p> </div>
<b>SAP HANA Studio</b>	<ol style="list-style-type: none"> <li>1. Start the SAP HANA studio.</li> <li>2. In the SAP HANA studio, add the SAP HANA system.</li> <li>3. Open the context menu (right-mouse click) in the <i>Systems</i> view, and select <i>Add System</i>. For more information about adding a system, see <i>Add an SAP HANA System</i> in the <i>SAP HANA Administration Guide</i> in Related Information.</li> <li>4. In the SAP HANA studio, log on to the system.</li> <li>5. From the context menu of the selected system, select <b>Lifecycle Management</b> <b>Platform Lifecycle Management</b> <b>SAP HANA Platform Lifecycle Management</b>.</li> </ol>
<b>SAP HANA Cockpit</b>	<ol style="list-style-type: none"> <li>1. Enter the SAP HANA cockpit URL in your browser. The URL depends on whether you are connecting to a single-container system or to a database in a multiple-container system. A single-container system is accessed through the URL: <code>http://&lt;host_FQDN&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code> For more information about the URLs in multiple-container systems, see <i>Configure HTTP Access to Multitenant Database Containers</i> in the <i>SAP HANA Administration Guide</i> in Related Information. <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>FQDN = fully qualified domain name</p> </div> </li> <li>2. The <i>SAP HANA Platform Lifecycle Management</i> tiles are visible on the homepage of the SAP HANA cockpit. If they are not, you can add them from the <i>SAP HANA Platform Lifecycle Management</i> tile catalog. For more information, see <i>Customizing the Homepage of SAP HANA Cockpit</i>.</li> </ol>

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.

Field Name	Description
<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 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 \[page 70\]](#)

### 4.10.3 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 *Registering in the System Landscape Directory using sldreg* and the introductory topic in this guide *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 (for example in the SAP HANA Studio) by modifying the `nameserver.ini` 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.

## Procedure

1. In SAP HANA studio in the Administration editor, choose the *Configuration* tab.
2. Right-click the `nameserver.ini` file and choose *Add Section*.
3. Create a new section `sld`.
4. 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
enable	Activates or deactivates the SLD data supplier	true	Allowed values are true, false. Re-enabling this parameter triggers a new generation of <code>sldreg.xml</code> and sending to the SLD system.
enable_virtdb-home	Activates or deactivates <code>SAP_IdenticalDatabaseSystem</code> in dependence of the system replication state and <code>sldvirtdbhome</code> parameter setting.		Allowed values are true, false.

Key	Meaning	Default Value	Note
Interval	Specifies the frequency (in seconds) with which the <code>sldreg.xml</code> file is generated. If a newly-generated document is the same as the previous one, it is not sent to the SLD.	300	It does not make sense to enter small positive values or negative values.  If you enter 0 or a negative value, data is transferred to the SLD only once.  Enter a value without a "1000 separator" (for example, 1899, not 1,899 or 1.899), otherwise it is interpreted as 0.
force_interval	Specifies how often (in seconds) the <code>sldreg.xml</code> 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 <code>slddest.cfg</code>  This file is a parameter for the call to <code>sldreg</code> .	<code>/usr/sap/ &lt;sid&gt;/SYS/ global</code>	Example: <code>/usr/sap/MPW/SYS/global</code>
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/ &lt;sid&gt;/ HDB&lt;id&gt;/ &lt;currenthost &gt;/trace</code>	Example: <code>/usr/sap/LRH/HDB42/velberlcm1/trace</code>

## Results

After modifying the parameters the data transfer will perform accordingly.

### **i** 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`.

Some additional resources related to Solution Manager are listed here:

- SAP Note 2082466 - Known issues in SAP HANA Platform lifecycle management (hdblcm)
- A troubleshooting guide for SAP Solution Manager is available on the SAP Community platform
- Also available on the SAP Community platform is a System Landscape Directory overview document.

---

## Related Information

[SAP Note 1018839 Registering in the System Landscape Directory using sldreg](#)

[SAP Community \(Technical Operations\): Troubleshooting Guide for SAP Solution Manager](#)

[SAP Community: System Landscape Directory Overview](#)

[Solution Manager documentation for Landscape Management Database \(LMDB\)](#)

[Configuring, Working with and Administering System Landscape Directory \(SAP NetWeaver\)](#)

[Database Trace \(Basic, User-Specific, and End-to-End\) \[page 466\]](#)

[Configuring an SAP HANA System to Connect to the System Landscape Directory \(SLD\) \[page 451\]](#)

## 4.11 Getting Support

Several functions are available in SAP HANA studio and the SAP HANA cockpit for offline administration to allow you or SAP Support to analyze and diagnose problems with SAP HANA.

### 4.11.1 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.

Diagnosis files are stored at the following default location on the SAP HANA server:

```
/usr/sap/<SID>/HDB<instance>/<host>/trace.
```

In a system with multitenant database containers, the trace files of the system database are stored at the above location. 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

[Monitoring Disk Space \[page 261\]](#)

[Multitenant Database Containers \[page 15\]](#)

## 4.11.1.1 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.

### Procedure

In the Administration editor, choose the *Diagnosis Files* tab.

### Results

All available diagnosis files are displayed. For more information about working with diagnosis files, see *Options for Diagnosis File Handling (SAP HANA Studio)*.

### Related Information

[Options for Diagnosis File Handling \(SAP HANA Studio\) \[page 461\]](#)

## 4.11.1.1.1 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 <b>sqltrace</b> . You can also filter by host.
Filter files by tenant database	To see the diagnosis files for a particular tenant database in a multiple-container system, 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.

Option	Description
Display file contents	<p>To display the contents of a file, right-click it and choose <a href="#">Open</a>, or double-click the file.</p> <p>The <a href="#">Show Start of File</a> and <a href="#">Show End of File</a> 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="603 539 1394 949" style="background-color: #fff9c4; padding: 10px;"> <p><b>i Note</b></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> <p><b>→ Tip</b></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 <a href="#">Merge Diagnosis Files...</a></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 <a href="#">Export to CSV file</a>, 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 <a href="#">Merged Diagnosis Files</a> 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 <a href="#">Delete</a>.</p> <p>You delete trace files in the following ways:</p> <ul style="list-style-type: none"> <li>You can delete trace files in the same way as other diagnosis files by right-clicking them and choosing <a href="#">Delete</a>.</li> </ul> <div style="background-color: #fff9c4; padding: 10px; margin: 10px 0;"> <p><b>i Note</b></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"> <li>You can batch delete trace files, for example all the trace files of a specific service, by choosing <a href="#">Delete Trace Files...</a> and making the required selection.</li> </ul> <div style="background-color: #fff9c4; padding: 10px; margin: 10px 0;"> <p><b>i Note</b></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 <a href="#">Compress</a>. 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 <a href="#">Download</a>. You can select multiple files to download.</p>

## 4.11.1.2 View Diagnosis Files in SAP HANA Cockpit

To diagnosis and analyze errors while SAP HANA is stopped or cannot be reached due to performance problems, you can access diagnosis files in the SAP HANA cockpit for offline administration.

### 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 SAP HANA cockpit for offline administration.

You can do this the following ways:

Option	Description
Directly	<p>Enter the URL in your browser:</p> <pre>https://&lt;host&gt;:1129/lmsl/hdbcockpit/&lt;sid&gt;/index.html</pre> <p><b>i Note</b></p> <p>It's also possible to access the SAP HANA cockpit for offline administration via the URL <code>http://&lt;host&gt;:1128/lmsl/hdbcockpit/&lt;sid&gt;/index.html</code>). However, this is not recommended because passwords are transferred in plain text via HTTP.</p>
From the SAP HANA cockpit	<ol style="list-style-type: none"><li>1. Open the SAP HANA cockpit by entering the URL in your browser: <code>https://&lt;host&gt;:43&lt;instance&gt;/sap/hana/admin/cockpit</code> (recommended) or <code>http://&lt;host&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code></li><li>2. In the <i>SAP HANA Database Administration</i> group, click the tile <i>SAP HANA Cockpit for Offline Administration</i>.</li></ol> <p><b>i Note</b></p> <p>If you access the SAP HANA cockpit via HTTP, then the SAP HANA cockpit for offline administration is also accessed via HTTP. Therefore, we recommend configuring the SAP HANA cockpit for HTTPS access.</p>

2. Open the *Diagnosis Files* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.

### **i Note**

You may be required to re-enter the password of the operating system user `<sid>adm`.

## Results

All available diagnosis files are displayed. Click a file to navigate to its contents. Alternatively, you can click the download button to download the file to the download directory of your browser on your client.

You can also search diagnosis files to find specific words and phrases.

### **i Note**

Text entered in the search field is treated as a regular expression. For example, a bracketed expression such as `[abc]` will find the single characters a, b and/or c. If you want to find square brackets (`[,]`), you must

escape them in the search expression using backslash (\). To find the exact sequence "abc", enter `abc` as the search expression.

## 4.11.2 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

### 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.

#### Note

To restore the default status or configuration of a trace, in the [Trace Configuration](#) dialog box choose [Restore Defaults](#).

## Results

Data recorded by traces is saved to trace files, which you can view on the *Diagnosis Files* tab.

## Related Information

[Database Trace \(Basic, User-Specific, and End-to-End\) \[page 466\]](#)

[SQL Trace \[page 469\]](#)

[Performance Trace \[page 474\]](#)

[Expensive Statements Trace \[page 475\]](#)

[Kernel Profiler \[page 477\]](#)

### 4.11.2.1 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 is written to the file:

```
<service>_alert_<host>.trc
```

You can access database trace files on the *Diagnosis Files* tab of the Administration editor.

#### Note

The trace files generated for the Web Dispatcher service are different. For more information, see *Trace Configuration for Internal Web Dispatcher*.

## Basic Database Trace Configuration

You configure the basic database trace in the Administration editor on the *Trace Configuration* tab.

Database tracing is always active. Information about error situations is always recorded.

If a trace component is available in all services, it can be configured for all services at once. You can also 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 at the ALL SERVICES level. Some components are only available in a particular service and cannot be changed at the ALL SERVICES level.

In the *Trace Configuration* dialog box, a trace level that has been inherited from the ALL SERVICES configuration (either the default or system configuration) is shown in brackets.

Not all trace components are visible by default in the *Trace Configuration* dialog box. To view all additional components, select *Show All Components*.

### 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.

The following trace levels are available:

- NONE (0)
- FATAL (1)
- ERROR (2)
- WARNING (3)
- INFO (4)
- DEBUG (5)

The higher the trace level, the more detailed the information recorded by the trace.

### Note

Even if you select NONE, information about error situations is still recorded.

## 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.

In the *Trace Configuration* dialog box, a trace level that has been inherited (either from the configured database trace or from the ALL SERVICES configuration in the user-specific trace) is shown in brackets.

### 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

### Note

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

[Traces and Trace Configuration for Internal Web Dispatcher \[page 478\]](#)

[End-to-End Analysis Overview](#)

## 4.11.2.2 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.

You configure the database trace for a tenant database in the Administration editor on the *Trace Configuration* tab.

The trace level of trace components 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 value in the *Database Trace Level* column.

### Note

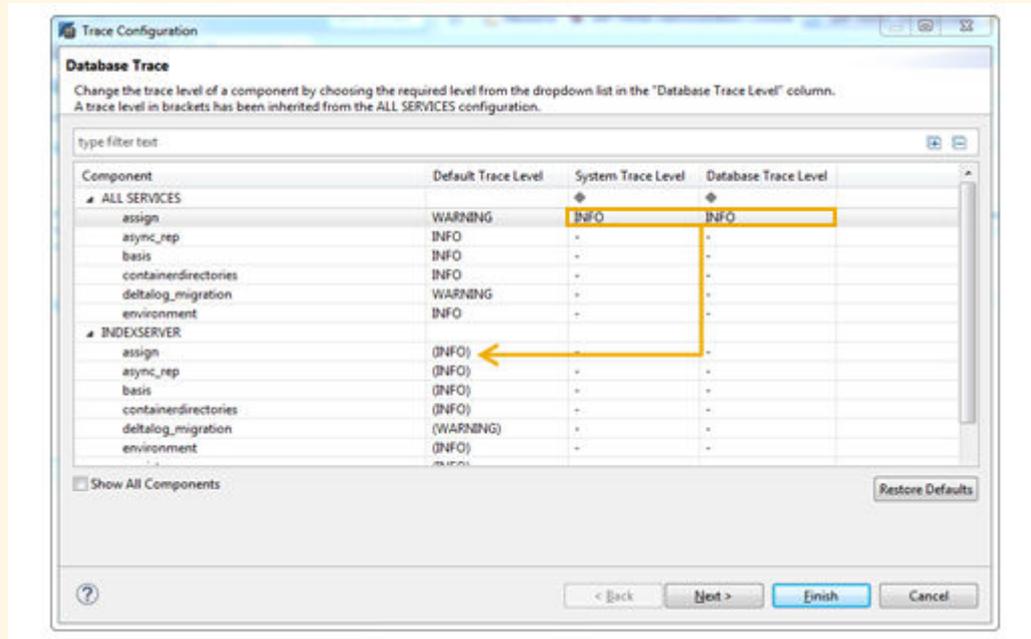
The database trace level of a component will also be displayed as the system trace level, and you cannot change the system trace level. This is because from the perspective of the tenant database, the database and the system are effectively the same. The true system trace level (that is, the value configured in the system database) is the default trace level for the tenant database.

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

### 4.11.2.3 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).

### Note

For more information on enabling the SQL trace in SAP HANA Studio, see KBA article [2031647](#) (How to enable SQL Trace in SAP HANA Studio).

### 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.

Trace Level	Description
ALL_WITH_RESULTS	<p>In addition to the trace generated with trace level ALL, the result returned by select statements is also included in the trace file.</p> <p><b>i Note</b> An SQL trace that includes results can quickly become very large.</p> <p><b>⚠ Caution</b> An SQL trace that includes results may expose security-relevant data, for example, query result sets.</p>

### Additional Configuration Options

Option	Configuration Parameter	Default	Description
Trace file name	tracefile	sqltrace_ \$HOST_ {PORT}_ {COUNT: 3}.py	<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>DB_&lt;dbname&gt;/sqltrace_&lt;HOST&gt;_ {PORT}_&lt;COUNT:3&gt;.py, where:</p> <ul style="list-style-type: none"> <li>DB_&lt;dbname&gt; is the sub-directory where the trace file is written if you are running on a tenant database</li> <li>\$HOST is the host name of the service (for example, indexserver)</li> <li>\$PORT is the port number of the service</li> <li>\$COUNT:3 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.</li> </ul>
User, application, object, and statement filters	user application_user application object statement_type	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 tables.</p> <p>All statements matching the filter criteria are recorded and saved to the specified trace file.</p>
Flush limit	flush_interval	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

To prevent SQL trace files from growing indefinitely, you can limit the size and number of trace files using the following parameters:

Parameter	Default	Description
<code>max_files</code>	<b>1</b>	Sets the maximum number of trace files  When the maximum number of trace files reached, the oldest trace file is deleted and a new one opened.  When set to 0, the trace file rotation is disabled.
<code>filesize_limit</code>	<b>1610612736</b> (or 1.5 GB)	Sets the maximum size of an individual trace file in bytes  When a trace file reaches the specified maximum file size, it is closed, and a new file created.  When set to 0, the file size is unlimited.

## 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.

### Note

For more information about the installation and usage of SAP HANA SQL Trace Analyzer, see KBA article [2412519](#)  (FAQ: SAP HANA SQL Trace Analyzer.)

## Related Information

[Diagnosis Files \[page 460\]](#)

[SAP Note 2036111](#) 

## 4.11.2.4 SQL Trace Options

Several options are available for configuring the SQL trace.

### Trace Levels

You can configure the following trace levels:

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	<p>In addition to the trace generated with trace level ALL, the result returned by select statements is also included in the trace file.</p> <p><b>i Note</b> An SQL trace that includes results can quickly become very large.</p> <p><b>⚠ Caution</b> An SQL trace that includes results may expose security-relevant data, for example, query result sets.</p>

## Other Configuration Options

The following additional configuration options are available:

Option	Description
Trace file	<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>DB_&lt;dbname&gt;/sqltrace_ \$HOST_ \${PORT}_ \${COUNT:3} .py, where:</p> <ul style="list-style-type: none"><li>• DB_&lt;dbname&gt; is the sub-directory where the trace file is written if you are running on a tenant database</li><li>• \$HOST is the host name of the service (for example, indexserver)</li><li>• \$PORT is the port number of the service</li><li>• \$COUNT:3 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.</li></ul>
User, application, object, and statement filters	<p>Filters to restrict traced statements to those of particular database or application users and applications, as well as to certain statement types and tables.</p> <p>All statements matching the filter criteria are recorded and saved to the specified trace file.</p>
Flush limit	<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>

### 4.11.2.5 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. It is inactive by default.

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).

It is recommended that you start performance tracing immediately before running the command(s) that you want to analyze and stop it immediately after they have finished. When you stop tracing, the results are saved to trace files that you can access on the [Diagnosis Files](#) tab of the Administration editor. You cannot analyze these files meaningfully in the SAP HANA studio, but instead must use a tool capable of reading the output format (\*.tpt). SAP Support has tools for evaluating performance traces.

You activate and configure the performance trace in the Administration editor on the [Trace Configuration](#) tab.

## Related Information

[Performance Trace Options \[page 475\]](#)

## 4.11.2.6 Performance Trace Options

Several options are available for configuring the performance trace.

### Standard Trace Mode

Option	Description
Trace file	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.

## 4.11.2.7 Expensive Statements Trace

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. It is inactive by default.

You activate and configure the expensive statements trace in the Administration editor on either the [Trace Configuration](#) tab or the [Performance](#) tab. Information about recorded expensive statements is displayed on the [Performance](#) tab.

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.

### Related Information

[Expensive Statements Trace Options \[page 476\]](#)

[Setting a Memory Limit for SQL Statements \[page 277\]](#)

## 4.11.2.8 Expensive Statements Trace Options

Several options are available for configuring the expensive statements trace.

Option	Description
Threshold duration	Threshold execution time in microseconds (default 1,000,000)
User, application filters, and object filters	Filters to restrict traced statements to those of particular database or application users, as well as to certain applications and tables
Trace parameter values	In SQL statements, field values may be specified as parameters (using a "?" in the syntax). If these parameter values are not required, you can deselect the <i>Trace parameter values</i> checkbox to reduce the amount of data traced.

### Additional Parameters

You can configure the expensive statement trace further using the following properties in the `expensive_statement` section of `global.ini` configuration file:

Property	Description
<code>trace_flush_interval</code>	Number of records after which trace file is flushed
<code>threshold_cpu_time</code>	Threshold CPU time of statement execution in microseconds  <b>i Note</b> Resource tracking and CPU time tracking must also be enabled. You can do this by configuring the corresponding parameters in the <code>resource_tracking</code> section of the <code>global.ini</code> file.
<code>threshold_memory</code>	Threshold memory usage of statement execution in bytes  <b>i Note</b> Resource tracking and memory tracking must also be enabled. You can do this by configuring the corresponding parameters in the <code>resource_tracking</code> section of the <code>global.ini</code> file.

## 4.11.2.9 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.

### Caution

To be able to use the kernel profile, you must have the SAP\_INTERNAL\_HANA\_SUPPORT role. This role is intended only for SAP HANA development support.

The kernel profile collects, for example, information about frequent and/or expensive execution paths during query processing.

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. When you stop tracing, the results are saved to trace files that you can access on the *Diagnosis Files* tab of the Administration editor. You cannot analyze these files meaningfully in the SAP HANA studio, but instead must use a tool capable of reading the configured output format, that is KCacheGrind or DOT (default format).

You activate and configure the kernel profile in the Administration editor on the *Trace Configuration* tab.

## Related Information

[Kernel Profiler Options \[page 477\]](#)

## 4.11.2.10 Kernel Profiler Options

Several options are available for configuring the kernel profiler.

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.

Option	Description
Database user, application user	The database user and/or application user you want to profile
Use KCachegrind format to write trace files	Output format of trace files (configurable when you stop tracing)

## 4.11.2.11 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` and can be viewed on the [Diagnosis Files](#) tab of the Administration editor.

You can configure the developer trace in the following ways:

- Changing the database trace level for the `dev_webdisp` component of the `webdispatcher` service  
You can do this on the [Trace Configuration](#) tab of the Administrator editor. The default trace level is ERROR.
- Changing (or adding) the property `rdisp/trace` in the `[profile]` section of the `webdispatcher.ini` configuration file  
You can do this on the [Configuration](#) tab of the Administration editor.  
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 on the [Trace Configuration](#) tab of the Administrator editor.

## 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 \(Basic, User-Specific, and End-to-End\) \[page 466\]](#)

[rdisp/TRACE\\* Parameters](#)

[icm/HTTP/trace\\_info](#)

[icm/HTTP/logging\\_<xx>](#)

[Logging in the ICM and SAP Web Dispatcher](#)

## 4.11.3 Configure 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 system and for individual services.

### Prerequisites

You have the system privilege INIFILE ADMIN.

### Procedure

1. In the SAP HANA studio, open the Administration editor and choose the *Configuration* tab.
2. Depending on whether you are configuring trace file rotation for all system services or for an individual service, proceed as follows:

Option	Description
<b>All services</b>	<ol style="list-style-type: none"><li>1. Navigate to the <code>global.ini</code> file and expand the section <code>trace</code>.</li><li>2. Configure the <code>maxfiles</code> parameter by specifying the maximum number of trace files that may exist.</li><li>3. Configure the <code>maxfilesize</code> parameter by specifying in bytes the maximum size an individual trace file may reach.</li></ol> <p><b>Note</b></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>
<b>Individual service</b>	<ol style="list-style-type: none"><li>1. Navigate to the configuration file of the relevant service (for example, <code>indexserver.ini</code>) and expand the section <code>trace</code>. If there is no <code>trace</code> section, create one by right-clicking the file and choosing <i>Add Section</i>.</li><li>2. Configure the <code>maxfiles</code> parameter by specifying the maximum number of trace files that may exist.</li><li>3. Configure the <code>maxfilesize</code> parameter specifying the maximum size an individual trace file may reach in bytes.</li></ol> <p><b>Note</b></p> <p>If the two parameters do not exist in the <code>trace</code> section or if you created a new <code>trace</code> section, create the parameters by right-clicking the section and choosing <i>Add Parameter</i>.</p>

### 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**

[Configure Traces in SAP HANA Studio \[page 465\]](#)

[Configure Check Thresholds \[page 251\]](#)

## **4.11.4 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**

[Troubleshoot Unresponsive System in SAP HANA Studio \[page 481\]](#)

[Troubleshoot an Unresponsive System in SAP HANA Cockpit \[page 482\]](#)

### **4.11.4.1 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.

### Note

In systems that support multitenant database containers, 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 \[page 484\]](#)

### 4.11.4.2 Troubleshoot an Unresponsive System in SAP HANA Cockpit

When a system cannot be reached by SQL or is experiencing performance problems, you cannot use the SAP HANA cockpit to troubleshoot and/or resolve issues. However, with the SAP HANA cockpit for offline administration 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 SAP HANA cockpit for offline administration.

You can do this the following ways:

Option	Description
<b>Directly</b>	<p>Enter the URL in your browser:</p> <pre>https://&lt;host&gt;:1129/lmsl/hdbcockpit/&lt;sid&gt;/index.html</pre> <p><b>i Note</b></p> <p>It's also possible to access the SAP HANA cockpit for offline administration via the URL <code>http://&lt;host&gt;:1128/lmsl/hdbcockpit/&lt;sid&gt;/index.html</code>). However, this is not recommended because passwords are transferred in plain text via HTTP.</p>
<b>From the SAP HANA cockpit</b>	<ol style="list-style-type: none"><li>1. Open the SAP HANA cockpit by entering the URL in your browser: <code>https://&lt;host&gt;:43&lt;instance&gt;/sap/hana/admin/cockpit</code> (recommended) or <code>http://&lt;host&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code></li><li>2. In the <i>SAP HANA Database Administration</i> group, click the tile <i>SAP HANA Cockpit for Offline Administration</i>.</li></ol> <p><b>i Note</b></p> <p>If you access the SAP HANA cockpit via HTTP, then the SAP HANA cockpit for offline administration is also accessed via HTTP. Therefore, we recommend configuring the SAP HANA cockpit for HTTPS access.</p>

2. On the homepage of the SAP HANA cockpit, click the *Troubleshoot Unresponsive System* tile.

The system starts collecting information about all connections, transactions, blocked transactions, and threads in the system.

## Results

Once system information is available, the *Troubleshoot Unresponsive System* app opens and displays the collected information on several tab pages. If necessary, you can cancel individual connections and transactions, or even cancel all transactions.

## 4.11.5 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.

## 4.11.6 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 studio, the SAP HANA cockpit, and the command line.

### 4.11.6.1 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>

To...	You Need...
List diagnosis information	<p>SELECT privilege on the view SYS.FULL_SYSTEM_INFO_DUMPS</p> <p>In the system database of a multiple-container system, you also need SELECT on SYS_DATABASES.FULL_SYSTEM_INFO_DUMPS so that you can see diagnosis information collected from tenant databases.</p>
Download collected diagnosis information	EXECUTE privilege on the procedure SYS.FULL_SYSTEM_INFO_DUMP_RETRIEVE
Delete collected diagnosis information	EXECUTE privilege on the procedure SYS.FULL_SYSTEM_INFO_DUMP_DELETE

- 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
<b>Collect all diagnosis information</b>	<p>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>.</p> <p><b>i</b> 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 <b>not</b> 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>

Option	Description
<b>Create and collect one or multiple sets of runtime environment (RTE) dump files</b>	<p>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:</p> <ul style="list-style-type: none"> <li>○ The index server(s) from which RTE dump files are to be collected</li> <li>○ The number of RTE dump file sets to be collected (possible values are 1, 2, 3, 4, and 5)</li> <li>○ 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.</li> </ul>

### **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.
5. Optional: Delete any old collections that you no longer need by selecting them and choosing *Delete Collections*.

## Related Information

[Diagnosis Information Collected \[page 491\]](#)

### 4.11.6.2 Collect and Download Diagnosis Information in SAP HANA 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. The SAP HANA cockpit for offline administration uses the SAP Host Agent to download diagnosis information when SAP HANA is both online and offline.

## 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 SAP HANA cockpit for offline administration.

You can do this the following ways:

Option	Description
<b>Directly</b>	<p>Enter the URL in your browser:</p> <pre>https://&lt;host&gt;:1129/lmsl/hdbcockpit/&lt;sid&gt;/index.html</pre> <p><b>i Note</b></p> <p>It's also possible to access the SAP HANA cockpit for offline administration via the URL <code>http://&lt;host&gt;:1128/lmsl/hdbcockpit/&lt;sid&gt;/index.html</code>). However, this is not recommended because passwords are transferred in plain text via HTTP.</p>
<b>From the SAP HANA cockpit</b>	<ol style="list-style-type: none"><li>1. Open the SAP HANA cockpit by entering the URL in your browser: <code>https://&lt;host&gt;:43&lt;instance&gt;/sap/hana/admin/cockpit</code> (recommended) or <code>http://&lt;host&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code></li><li>2. In the <i>SAP HANA Database Administration</i> group, click the tile <i>SAP HANA Cockpit for Offline Administration</i>.</li></ol> <p><b>i Note</b></p> <p>If you access the SAP HANA cockpit via HTTP, then the SAP HANA cockpit for offline administration is also accessed via HTTP. Therefore, we recommend configuring the SAP HANA cockpit for HTTPS access.</p>

2. Open the *Diagnosis Files* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.
3. Open the *Diagnosis Information Collections* tab.  
Any existing collections of diagnosis information are listed.
4. In the footer bar, first click *Collect*, then specify the scope of information to be collected by clicking one of the following options:

Option	Description
<b>Diagnosis Information</b>	Select this option if you want to collect all diagnosis information for a specific time period, by default the last 7 days.

Option	Description
<b>RTE Dump Files</b>	<p>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:</p> <ul style="list-style-type: none"> <li>○ The hosts from which RTE dump files are to be collected</li> <li>○ The number of RTE dump file sets to be collected (possible values are 1, 2, 3, 4, and 5)</li> <li>○ 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.</li> </ul>

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. For more information about the information collected, see *Diagnosis Information Collected*.

## Related Information

[Diagnosis Information Collected \[page 491\]](#)

### 4.11.6.3 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 since tenant database administrators do not have operating system access. Tenant database

administrators must use the SAP HANA studio to collect diagnosis information from their database (see *Collect and Download Diagnosis Information in SAP HANA Studio*).

## 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:

- Single-container system: `fullsysteminfodump_<HOST>_<SID>_<timestamp>.zip`
- Multiple-container system: `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, but you can look it up with the command: `hdbstrvutil -z | grep DIR_GLOBAL=`

## Related Information

[Collect and Download Diagnosis Information in SAP HANA Studio \[page 484\]](#)

[fullSystemInfoDump.py Command Line Options \[page 490\]](#)

[Diagnosis Information Collected \[page 491\]](#)

## 4.11.6.4 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  <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>If you are connected to the system database of a multi-container system, only information from the system views of the system database will be collected. Information from the system views of tenant databases will <b>not</b> be collected regardless of this option.</p> </div>
<code>--file &lt;filename&gt;</code>	Zips the specified file in its source directory  <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>Note: This option only zips the file; it does not trigger the collection of any other information.</p> </div>
<code>--days &lt;no. of days&gt;</code>	Collects information from the specified number of past days The default value is 7.  <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>You cannot use this option with the options <code>--fromDate</code> and <code>--toDate</code>.</p> </div>
<code>--fromDate &lt;YYYY-MM-DD&gt;</code>	Collects information starting from the specified date
<code>--toDate &lt;YYYY-MM-DD&gt;</code>	Collects information up to the specified date
<code>--rtedump</code>	Restricts the information collected to an RTE dump file or files  You can configure the creation and collection of RTE dump files further with the remaining options.
<code>--indexservers &lt;comma-separated list of index servers&gt;</code>	Specifies the index server(s) from which RTE dump files are to be collected  By default, dump files are created and collected for all index servers
<code>--interval &lt;interval in minutes&gt;</code>	Specifies the interval at which RTE dump files are to be collected  Possible values are 1, 5, 10, 15, and 30. The default value is 1.

Option	Description
<code>--sets &lt;no. of RTE dump file sets&gt;</code>	Specifies the number of RTE dump file sets to be collected. Possible values are 1, 2, 3, 4, and 5.
<code>--tenant &lt;database name in multiple-container system&gt;</code>	Specifies which database in a multiple-container system information is to be collected from You must specify a database name. To collect information from the system database, specify <b>SYSTEMDB</b> .  <b>i Note</b> Do not use this option if executing the script in a single-container system.

## 4.11.6.5 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 studio, the SAP HANA cockpit for offline administration, or directly from the command line.

### **i Note**

All of the following file types are collected unless the option `--rtdump` 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/mimetypemapping.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.

### **i** Note

If you trigger the collection of diagnosis information from the SAP HANA cockpit for offline administration, information from system views cannot be collected since it does not use an SQL connection.

- 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']
```

---

## 4.11.7 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.

### 4.11.7.1 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 attach to a support message for example.

#### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Used Memory](#), [Disk Usage](#), or [CPU Usage](#) tiles are visible on the homepage of the SAP HANA cockpit. If they're not, you can add them again from the tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

#### Procedure

1. Open the [Performance Monitor](#) app by clicking the [CPU Usage](#), [Disk Usage](#), or [Used Memory](#) tile on the homepage of the SAP HANA cockpit. The [Performance Monitor](#) app opens displaying the load graph for the selected resource: CPU, disk, or memory. For more information, see [Analyze Past Performance](#) in Related Information.
2. Select [Export All](#) in the footer bar to export the CPU, disk, and memory KPI data as a single performance monitor data set. The data set includes data collected within the time frame of the past 12 days.

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

[Monitor and Analyze Past Performance \[page 313\]](#)

[Import Performance Monitor Data \[page 497\]](#)

---

## 4.11.7.2 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 SAP HANA cockpit.

### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`. You can grant roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [Support Tools](#) tile is visible on the homepage of the SAP HANA cockpit. If it is not, you can add it from the tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).
- You have access to a zip file, containing a performance monitor data set of an SAP HANA system, which was created using the [Export All](#) feature in the [Used Memory](#), [Disk Usage](#), or [CPU Usage](#) apps.

### Procedure

1. Open the [Support Tools](#) app by clicking the [Support Tools](#) tile on the homepage of the SAP HANA cockpit. The [Support Tools](#) app opens displaying a list of imported [Performance Monitor Data Sets](#).
2. Select [Import](#) and select the file containing the performance monitor data set that you want to import.
3. Enter a description and select [Import](#).

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](#) app opens and displays the KPI data stored inside the data set. Use the [Performance Monitor](#) app to visually analyze historical performance data across a range of key performance indicators related in particular to memory, disk, and CPU usage.

### Related Information

[Monitor and Analyze Past Performance \[page 313\]](#)

[Export Performance Monitor Data \[page 496\]](#)

---

## 4.11.8 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 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

[SAP Note 1634848](#)

[SAP Note 37001](#)

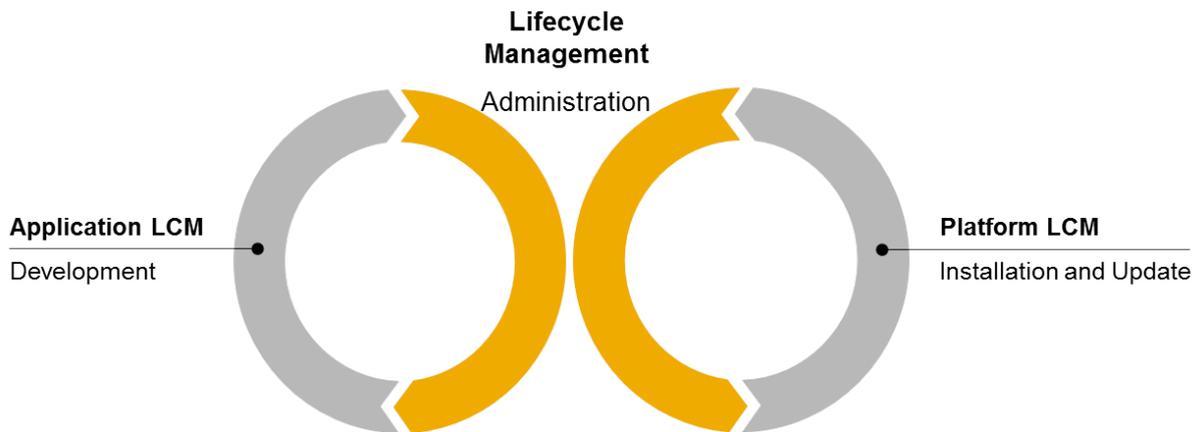
[SAP Note 1592925](#)

[SAP Note 1058533](#)

[SAProuter](#)

## 5 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. For example, by renaming your SAP HANA system, relocating the system to different hardware, or converting the system to a multiple-container enabled 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

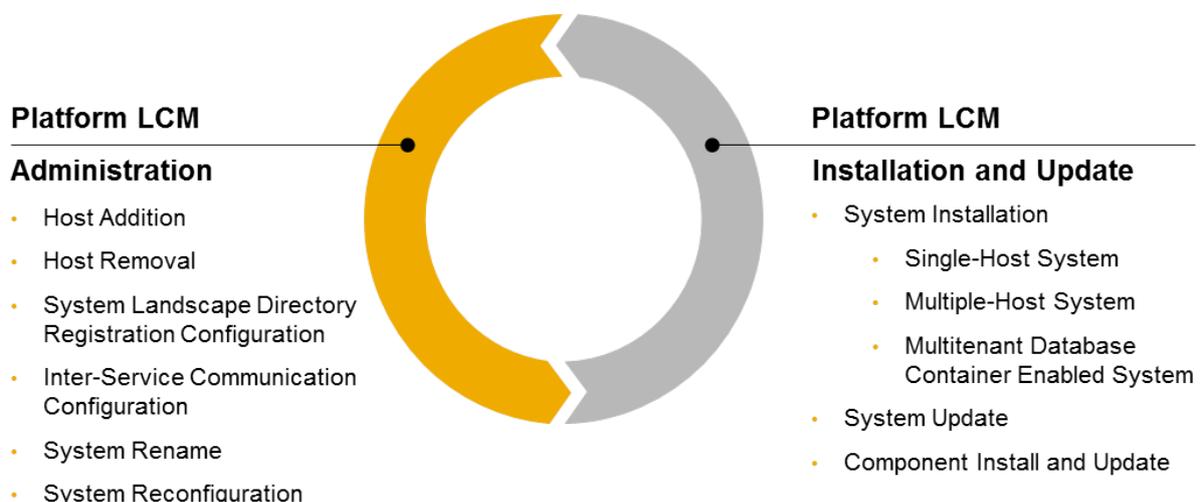
[SAP HANA Platform Lifecycle Management \[page 500\]](#)

[SAP HANA Application Lifecycle Management \[page 599\]](#)

## 5.1 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 \[page 502\]](#)

[Configuring a Multiple-Host System \[page 532\]](#)

[Configuring an SAP HANA System to Connect to the System Landscape Directory \(SLD\) \[page 451\]](#)

[Configuring SAP HANA Inter-Service Communication \[page 565\]](#)

[Rename an SAP HANA System Host \[page 585\]](#)

[Change the SID of an SAP HANA System \[page 587\]](#)

[Change the Instance Number of an SAP HANA System \[page 589\]](#)

[Relocate the SAP HANA System \[page 591\]](#)

[Copy or Clone an SAP HANA System \[page 594\]](#)

[Converting an SAP HANA System to Support Multitenant Database Containers \[page 572\]](#)

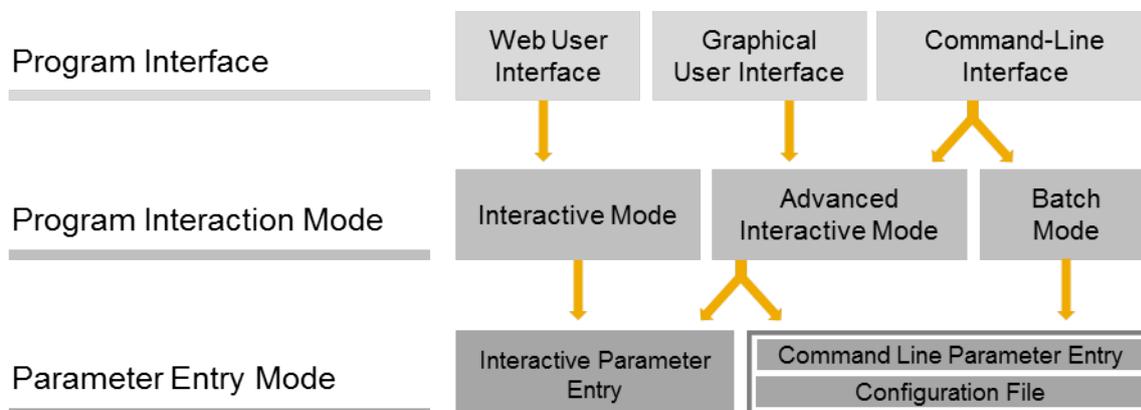
## 5.1.1 About the SAP HANA Database Lifecycle Manager (HDBLCM)

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.

### 5.1.1.1 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.



- [Performing LCM Tasks by Program Interface \[page 504\]](#)
- [Use the Web User Interface to Perform Platform LCM Tasks \[page 509\]](#)
- [Use the Graphical User Interface to Perform Platform LCM Tasks \[page 504\]](#)
- [Use the Command-Line Interface to Perform Platform LCM Tasks \[page 505\]](#)
- [Performing LCM Tasks by Program Interaction Mode \[page 512\]](#)
- [Use Interactive Mode to Perform Platform LCM Tasks \[page 512\]](#)
- [Use Advanced Interactive Mode to Perform Platform LCM Tasks \[page 513\]](#)
- [Use Batch Mode to Perform Platform LCM Tasks \[page 515\]](#)
- [Performing LCM Tasks by Parameter Entry Method \[page 517\]](#)
- [Entering Platform LCM Parameters Interactively \[page 518\]](#)
- [Entering Platform LCM Parameters as Call Options from the Command Line \[page 520\]](#)

- [Use LCM Configuration Files to Enter Parameters \[page 518\]](#)

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 or from the SAP HANA studio (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.

### 5.1.1.1.1 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`.

## 5.1.1.1.2 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 \[page 504\]](#)

[Use the Command-Line Interface to Perform Platform LCM Tasks \[page 505\]](#)

[Using the Web User Interface \[page 507\]](#)

### 5.1.1.1.2.1 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 &lt;installation medium&gt;/DATA_UNITS/ HDB_LCM_LINUX_X86_64</pre>
Installation Medium (IBM Power Systems)	<pre>cd &lt;installation medium&gt;/DATA_UNITS/ HDB_LCM_LINUX_PPC64</pre>
SAP HANA resident HDBLCM	<pre>cd &lt;sapmnt&gt;/&lt;SID&gt;/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 \[page 503\]](#)

[Entering Platform LCM Parameters as Call Options from the Command Line \[page 520\]](#)

## 5.1.1.1.2 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
<b>Installation Medium (Intel-Based Hardware Platforms)</b>	<pre>cd &lt;installation medium&gt;/DATA_UNITS/ HDB_LCM_LINUX_X86_64</pre>
<b>Installation Medium (IBM Power Systems)</b>	<pre>cd &lt;installation medium&gt;/DATA_UNITS/ HDB_LCM_LINUX_PPC64</pre>
<b>SAP HANA resident HDBLCM</b>	<pre>cd &lt;sapmnt&gt;/&lt;SID&gt;/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:

```
./hdblcmm
```

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 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 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.

#### **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 \[page 503\]](#)

[Performing LCM Tasks by Parameter Entry Method \[page 517\]](#)

[Performing LCM Tasks by Program Interaction Mode \[page 512\]](#)

[Entering Platform LCM Parameters as Call Options from the Command Line \[page 520\]](#)

## 5.1.1.1.2.3 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.

### Related Information

[About the Web User Interface \[page 507\]](#)

[Use the Web User Interface to Perform Platform LCM Tasks \[page 509\]](#)

[Log Off From an SAP HANA System \[page 511\]](#)

[Troubleshooting the Web User Interface \[page 512\]](#)

### 5.1.1.1.2.3.1 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.

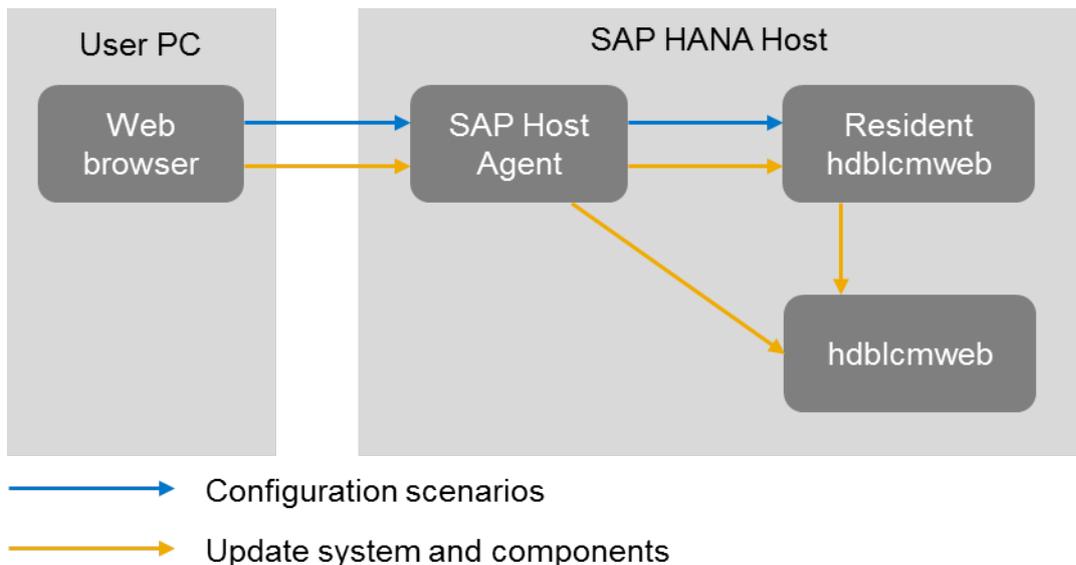
The backend is provided by the special executable `hdb1cmweb`, which is started automatically by the SAP Host Agent as soon as an action is triggered from the Web user interface and terminates after the action completes.

#### **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 \[page 523\]](#)

## 5.1.1.1.2.3.2 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 Platform Lifecycle Management view within the SAP HANA studio.

### 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](#).
    - 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* in Related Information.

- You are logged on as the system administrator user `<sid>adm`.

You should verify that the following additional prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from the SAP HANA studio.

- The SAP HANA studio revision is 120 or higher.
- For Linux:
  - The system property `org.eclipse.swt.browser.XULRunnerPath` should be set in `hdbstudio.ini` to point to the path of XULRunner, for example:  
`-Dorg.eclipse.swt.browser.XULRunnerPath=<path to xulrunner>`.  
This `hdbstudio.ini` file is located in the same folder as the executable that is used to start the SAP HANA studio. For Linux, the default location is `hana/shared/<SID>/hdbstudio..`

## 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
<b>Web Browser</b>	<p>Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser:</p> <p><code>https://&lt;hostname&gt;:1129/lmsl/HDBLCM/&lt;SID&gt;/index.html</code></p> <div style="background-color: #fff9c4; padding: 5px;"><p><b>i Note</b></p><p>The URL is case sensitive. Make sure you enter upper and lower case letters correctly.</p></div>
<b>SAP HANA Studio</b>	<ol style="list-style-type: none"><li>1. Start the SAP HANA studio.</li><li>2. In the SAP HANA studio, add the SAP HANA system.</li><li>3. Open the context menu (right-mouse click) in the <a href="#">Systems</a> view, and select <a href="#">Add System</a>. For more information about adding a system, see <a href="#">Add an SAP HANA System</a> in the <a href="#">SAP HANA Administration Guide</a> in Related Information.</li><li>4. In the SAP HANA studio, log on to the system.</li><li>5. From the context menu of the selected system, select <a href="#">Lifecycle Management</a> &gt; <a href="#">Platform Lifecycle Management</a> &gt; <a href="#">SAP HANA Platform Lifecycle Management</a>.</li></ol>
<b>SAP HANA Cockpit</b>	<ol style="list-style-type: none"><li>1. Enter the SAP HANA cockpit URL in your browser. The URL depends on whether you are connecting to a single-container system or to a database in a multiple-container system. A single-container system is accessed through the URL: <code>http://&lt;host_FQDN&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code> For more information about the URLs in multiple-container systems, see <a href="#">Configure HTTP Access to Multitenant Database Containers</a> in the <a href="#">SAP HANA Administration Guide</a> in Related Information.</li></ol> <div style="background-color: #fff9c4; padding: 5px;"><p><b>i Note</b></p><p>FQDN = fully qualified domain name</p></div>

Option	Description
	2. The <i>SAP HANA Platform Lifecycle Management</i> tiles are visible on the homepage of the SAP HANA cockpit. If they are not, you can add them from the <i>SAP HANA Platform Lifecycle Management</i> tile catalog. For more information, see <i>Customizing the Homepage of SAP HANA Cockpit</i> .

## Results

The SAP HANA database lifecycle manager is displayed as a Web user interface in either a standalone browser or in the SAP HANA studio.

## Related Information

[SAPUI5 Developer Guide](#)

[Add an SAP HANA System \[page 70\]](#)

### 5.1.1.1.2.3.3 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.

## 5.1.1.1.2.3.4 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 - Troubleshooting note for SAP HANA Platform Management tool hdblcmm](#)

## 5.1.1.1.3 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.

### Related Information

[Use Interactive Mode to Perform Platform LCM Tasks \[page 512\]](#)

[Use Advanced Interactive Mode to Perform Platform LCM Tasks \[page 513\]](#)

[Use Batch Mode to Perform Platform LCM Tasks \[page 515\]](#)

### 5.1.1.1.3.1 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

To access the SAP HANA database lifecycle manager Web user interface, see Related Information.

## Procedure

1. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	<code>cd &lt;installation medium&gt;/DATA_UNITS/HDB_LCM_LINUX_X86_64</code>
Installation Medium (IBM Power Systems)	<code>cd &lt;installation medium&gt;/DATA_UNITS/HDB_LCM_LINUX_PPC64</code>
SAP HANA resident HDBLCM	<code>cd &lt;sapmnt&gt;/&lt;SID&gt;/hdblcmm</code>

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:

Option	Description
Graphical Interface	<code>./hdblcmmgui</code>
Command-line Interface	<code>./hdblcmm</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 parameter 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 \[page 503\]](#)

[Use the Web User Interface to Perform Platform LCM Tasks \[page 509\]](#)

### 5.1.1.1.3.2 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 perform platform LCM tasks in interactive mode, but 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

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
<b>Installation Medium (Intel-Based Hardware Platforms)</b>	<pre>cd &lt;installation medium&gt;/DATA_UNITS/ HDB_LCM_LINUX_X86_64</pre>
<b>Installation Medium (IBM Power Systems)</b>	<pre>cd &lt;installation medium&gt;/DATA_UNITS/ HDB_LCM_LINUX_PPC64</pre>
<b>SAP HANA resident HDBLCM</b>	<pre>cd &lt;sapmnt&gt;/&lt;SID&gt;/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.

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 \[page 503\]](#)

### 5.1.1.1.3.3 Use Batch Mode to Perform Platform LCM Tasks

Batch mode is a method for running SAP HANA platform lifecycle management (LCM) tools which starts the program and runs it to completion without requiring you to interact with it any further. Batch mode must be run with the SAP HANA platform LCM command-line tools. 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

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)	<code>cd &lt;installation medium&gt;/DATA_UNITS/HDB_LCM_LINUX_X86_64</code>
Installation Medium (IBM Power Systems)	<code>cd &lt;installation medium&gt;/DATA_UNITS/HDB_LCM_LINUX_PPC64</code>
SAP HANA resident HDBLCM	<code>cd &lt;sapmnt&gt;/&lt;SID&gt;/hdblcm</code>

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 --batch <additional parameters>
```

or

```
./hdblcm -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

[Use LCM Configuration Files to Enter Parameters \[page 518\]](#)

[Entering Platform LCM Parameters as Call Options from the Command Line \[page 520\]](#)

### 5.1.1.1.4 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:

<b>Interactively (Default)</b>	Using either command line interface or graphical 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.
<b>Command Line Options</b>	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 ( <code>--help</code> ) for the individual SAP HANA lifecycle management tool.
<b>Configuration File</b>	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

## Related Information

[Entering Platform LCM Parameters Interactively \[page 518\]](#)

[Use LCM Configuration Files to Enter Parameters \[page 518\]](#)

[Entering Platform LCM Parameters as Call Options from the Command Line \[page 520\]](#)

## 5.1.1.1.4.1 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.

## 5.1.1.1.4.2 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
<b>Installation Medium (Intel-Based Hardware Platforms)</b>	<pre>cd &lt;installation medium&gt;/DATA_UNITS/ HDB_LCM_LINUX_X86_64</pre>

Option	Description
<b>Installation Medium (IBM Power Systems)</b>	<code>cd &lt;installation medium&gt;/DATA_UNITS/ HDB_LCM_LINUX_PPC64</code>
<b>SAP HANA resident HDBLCM</b>	<code>cd &lt;sapmnt&gt;/&lt;SID&gt;/hdbldcm</code>

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.

```
./hdbldcm --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.

```
./hdbldcm --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 \[page 503\]](#)

---

### 5.1.1.1.4.3 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.

### 5.1.1.1.5 Executing Platform LCM Tasks

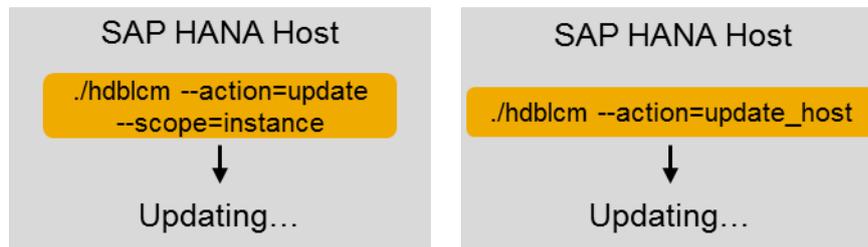
SAP HANA platform lifecycle management tasks can be performed on multiple-host systems centrally, by running the SAP HANA database lifecycle manager (HDBLCM) from any worker host and using remote execution to replicate the call on all remaining system hosts. Otherwise, the platform LCM tasks can be executed 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 \[page 521\]](#)

[Decentralized Execution of Platform LCM Tasks \[page 526\]](#)

### 5.1.1.1.5.1 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.

#### 5.1.1.1.5.1.1 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 install the SAP Host Agent on all system hosts. In order to use SSH, the SFTP subsystem

must be active. 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

[1944799 - SAP HANA Guidelines for SLES Operating System Installation](#)

[2009879 - SAP HANA Guidelines for Red Hat Enterprise Linux \(RHEL\) Operating System](#)

### 5.1.1.1.5.1.2 Using SAP Host Agent to Execute Platform LCM Tasks

In previous SAP HANA releases, it was only possible to perform multiple-host system tasks by providing root credentials and executing platform on remote hosts via secure shell (SSH). Since SAP HANA Support Package Stack (SPS) 09, it has been possible to perform platform LCM tasks without root credentials by using the SAP Host Agent.

Even though the SAP Host Agent is not required to be installed on the SAP HANA system, the SAP HANA database lifecycle manager (HDBLCM) heavily relies on it 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 for offline administration also 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. For more information, see *SAP HANA Cockpit for Offline Administration*.

The SAP Host Agent is installed and updated by default during SAP HANA system installation and update, unless the call option `--install_hostagent=off` is used. We recommend installing and updating the SAP Host Agent with the SAP HANA server to ensure version compatibility, however in some cases you may need to install or update only the SAP Host Agent. For information about installing or updating the SAP Host Agent individually, see *Installing SAP Host Agent Manually* and *Upgrading SAP Host Agent Manually* in Related Information.

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 addition of new host to an SAP HANA system (also during 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

[SAP HANA Cockpit for Offline Administration \[page 56\]](#)

[Installing SAP Host Agent Manually](#)

[Updating SAP Host Agent Manually](#)

[SSL Configuration for the SAP Host Agent](#)

### 5.1.1.1.5.1.2.1 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 and the SAP HANA cockpit for offline administration 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`.
- 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](#)

### 5.1.1.1.5.1.2.2 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

#### **i** Note

The SAP HANA cockpit for offline administration also 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. For more information, see *SAP HANA Cockpit for Offline Administration*.

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).

## Related Information

[SAP HANA Cockpit for Offline Administration \[page 56\]](#)

## 5.1.1.1.5.1.2.3 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* in Related Information.
- 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* in Related Information.
- 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 --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. If your system is SAP HANA multitenant database container (multiple-container) enabled system, 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* in Related Information.

## Related Information

[Host Addition Concepts \[page 535\]](#)

[Using the SAP HANA Platform LCM Tools \[page 502\]](#)

[Using SAP Host Agent to Execute Platform LCM Tasks \[page 522\]](#)

[Installing SAP Host Agent Manually](#)

[Updating SAP Host Agent Manually](#)

[Configuring SSL for SAP Host Agent on UNIX](#)

### 5.1.1.1.5.2 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 - Performing SAP HANA platform lifecycle management administration tasks on multiple-host systems without SSH or root credentials](#)

[Executing Platform LCM Tasks \[page 520\]](#)

[Centralized Execution of Platform LCM Tasks \[page 521\]](#)

### 5.1.1.1.6 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 where log files and traces are stores, Linux kernel parameter settings, or using the underlying LCM tools for troubleshooting purposes.

## Related Information

[Logging \[page 527\]](#)

[Linux Kernel Parameters \[page 528\]](#)

[General Troubleshooting for the SAP HANA Platform LCM Tools \[page 529\]](#)

#### 5.1.1.1.6.1 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 more easy

You can also view the last three log files in the SAP HANA studio using the administration function *Diagnosis Files*. For more information, see 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`

## 5.1.1.1.6.2 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.

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 <code>ispd</code> ating 1 GB)	1073741824	<code>/etc/sysctl.conf</code>
<code>kernel.shmni</code>	Maximum number of shared memory segments	<ul style="list-style-type: none"> <li>RAM <math>\geq</math> 256 GB: 524288</li> <li>RAM <math>\geq</math> 64 GB: 65536</li> <li>RAM &lt; 64 GB: 4096</li> </ul>	<code>/etc/sysctl.conf</code>

Parameter	Description	Value	Location
kernel.shmall	System-wide limit of total shared memory, in 4k pages	<ul style="list-style-type: none"> <li>RAM &gt;= 35.5 TB: (shmmax * shmmni) / 65536</li> <li>RAM &lt; 35.5 TB: (0.9 * &lt;RAM in bytes&gt;) / 4096</li> </ul>	/etc/sysctl.conf
net.ipv4.ip_local_port_range	Lower limit of ephemeral port range	40000	/etc/sysctl.conf
vm.max_map_count	Maximum number of Virtual Memory Areas (VMAs) that a process can own	<ul style="list-style-type: none"> <li>Intel-Based Hardware Platforms: 1000000 + &lt;RAM in GB&gt; * 32768</li> <li>IBM Power Systems: 1000000 + &lt;RAM in GB&gt; * 16384</li> </ul> <p>Maximum value: 2147483647</p>	/etc/sysctl.conf

### 5.1.1.1.6.3 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

Program Name	Description	Location
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
hdbreg	Command-line tool for registering an SAP HANA system	<installation path>/<SID>/ global/hdb/install/bin and /usr/sap/<SID>/SYS/ global/hdb/install/bin
hdbremovehost	Command-line tool for removing a host	<installation path>/<SID>/ global/hdb/install/bin and /usr/sap/<SID>/SYS/ global/hdb/install/bin
hdbmodify	This command line tool removes and adds remote hosts.  Furthermore, the listen interface can be changed ('local', 'global', 'internal').	<installation path>/<SID>/ global/hdb/install/bin and /usr/sap/<SID>/SYS/ global/hdb/install/bin
hdbupprep	Command-line tool for upgrading a repository by loading delivery units into the database	<installation path>/<SID>/ global/hdb/install/bin and /usr/sap/<SID>/SYS/ global/hdb/install/bin

## 5.1.1.2 Users Created During Installation

The following users are automatically created during the installation: <sid>adm, sapadm, and SYSTEM.

User	Description
<sid>adm	<p>The operating system administrator.</p> <ul style="list-style-type: none"><li>• The user &lt;sid&gt;adm is the operating system user required for administrative tasks such as starting and stopping the system.</li><li>• The user ID of the &lt;sid&gt;adm 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.</li><li>• The password of the &lt;sid&gt;adm user is set during installation with the <code>password</code> parameter.</li></ul>
sapadm	<p>The SAP Host Agent administrator.</p> <ul style="list-style-type: none"><li>• 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>.</li><li>• 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.</li><li>• The password of the <code>sapadm</code> user is set during installation with the <code>sapadm_password</code> parameter.</li></ul>
SYSTEM	<p>The database superuser.</p> <ul style="list-style-type: none"><li>• 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.</li><li>• The password of the <code>SYSTEM</code> user is set during installation with the <code>system_user_password</code> parameter.</li></ul>

### Related Information

[Operating System User <sid>adm \[page 647\]](#)

## 5.1.2 Configuring the SAP HANA System

After installation the SAP HANA system can be configured for compatibility with other SAP products or reconfigured from the original installation settings.

### Related Information

---

[Configuring a Multiple-Host System \[page 532\]](#)

[Configuring Host Roles \[page 554\]](#)

[Configuring SAP HANA Inter-Service Communication \[page 565\]](#)

[Converting an SAP HANA System to Support Multitenant Database Containers \[page 572\]](#)

## 5.1.2.1 Configuring a Multiple-Host System

It is possible to add hosts after installation to a single-host or multiple-host SAP HANA system.

Before adding a host, it is important to review multiple-host system concepts, for example storage and networking, and also review the SAP HANA database lifecycle manager add host technology concepts, for example the differences between local and remote host addition or configuring the listen interface.

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*.

### 5.1.2.1.1 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

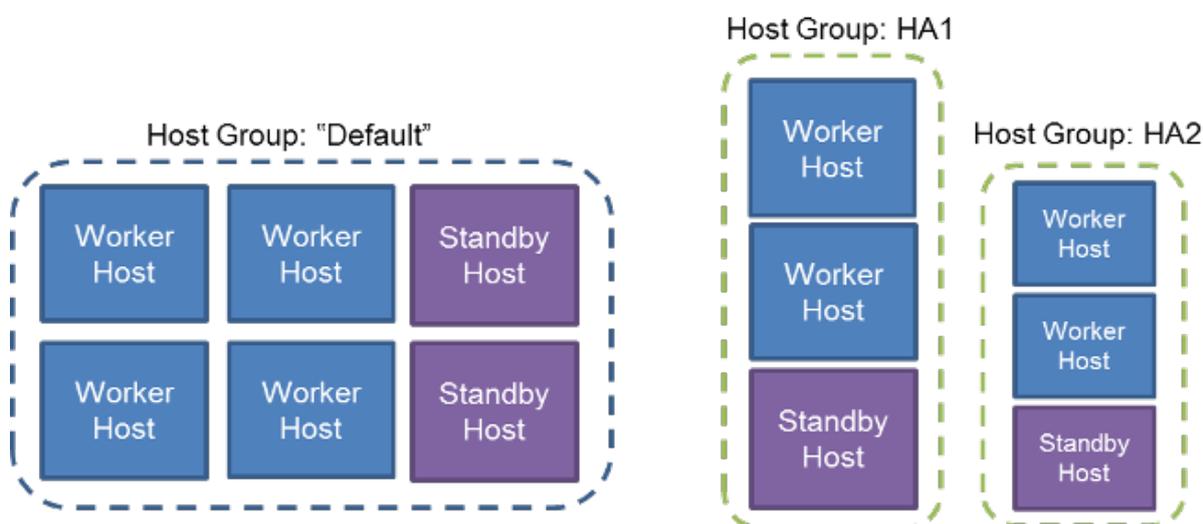
As an in-memory database, SAP HANA is not only concerned with maintaining the reliability of its data in the event of failures, but also with resuming operations with most of that data loaded back in memory as quickly as possible. Host auto-failover is a local fault recovery solution that can be used 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.

## 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:

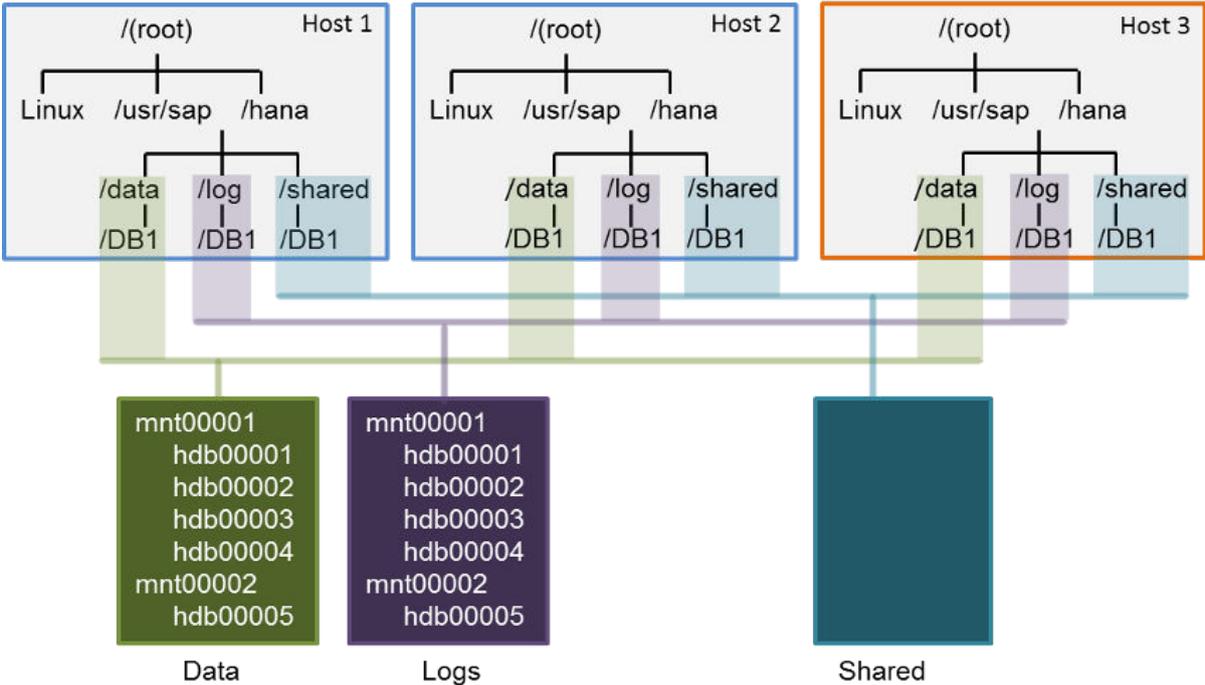
- The standby host has file access to data and log volumes of the failed host.
- The failed worker host no longer has access to write to files - called fencing.

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

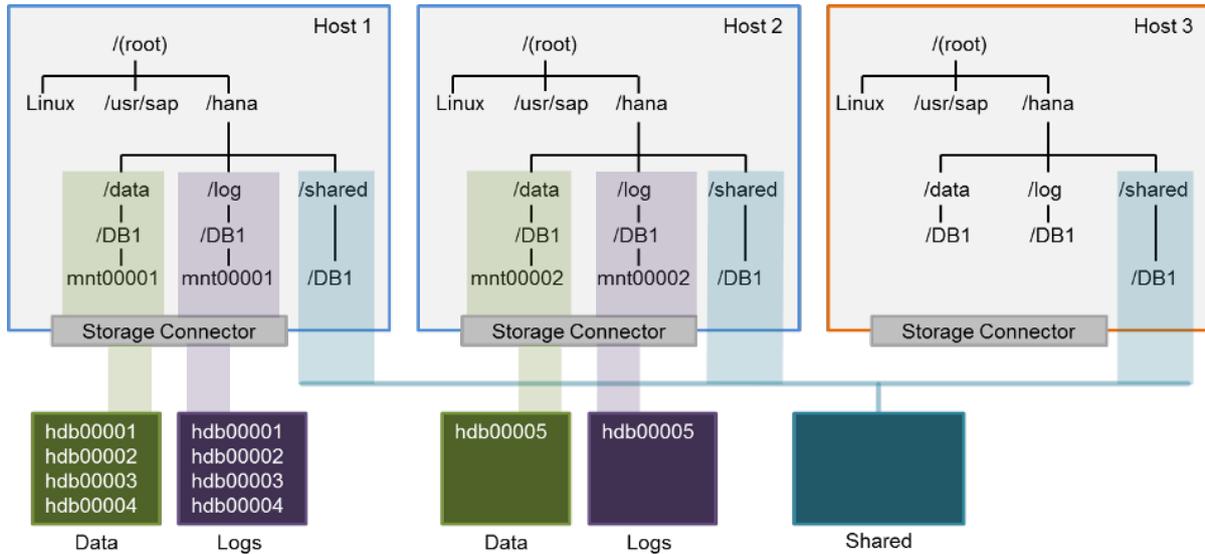
A shared storage subsystem, which is accessed using file systems such as NFS or IBM's GPFS, makes it easy to ensure that the standby host has access to all active host files in the system. In a shared storage solution, the externally attached storage subsystem devices are capable of providing dynamic 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.



### 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

[Setting Up Host Auto-Failover \[page 850\]](#)

[1900823 - SAP HANA Storage Connector API](#)

[405827 - Linux: Recommended file systems](#)

### 5.1.2.1.2 Host Addition Concepts

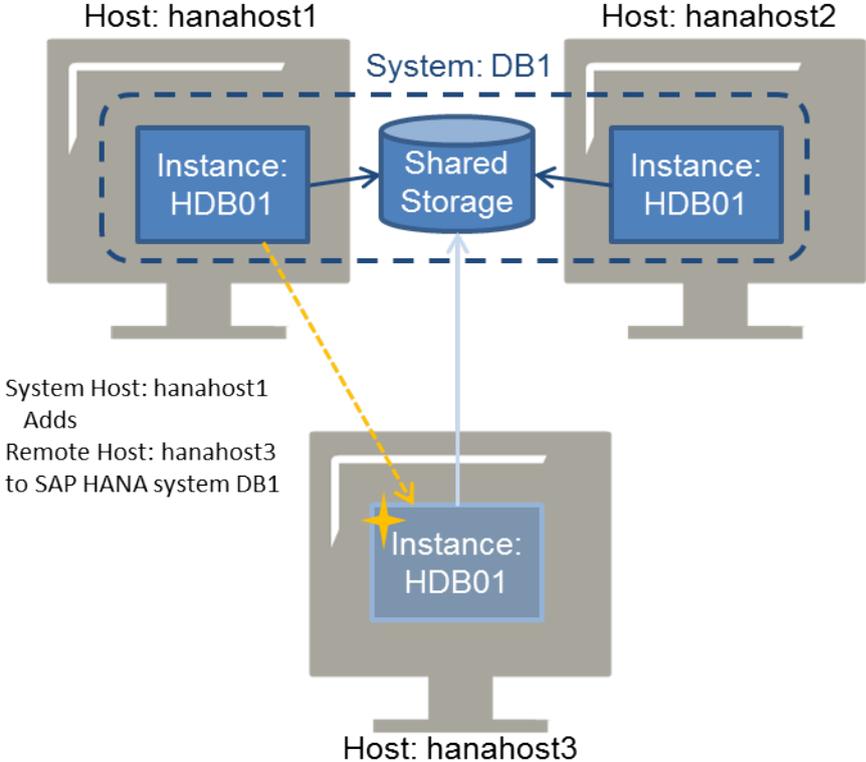
You can add hosts to an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM).

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.

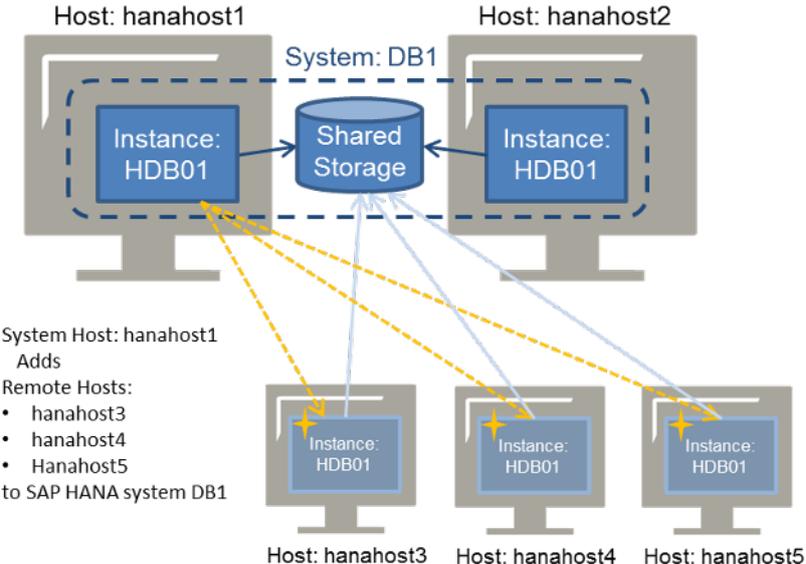
#### Adding Hosts from an Integrated Host

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.

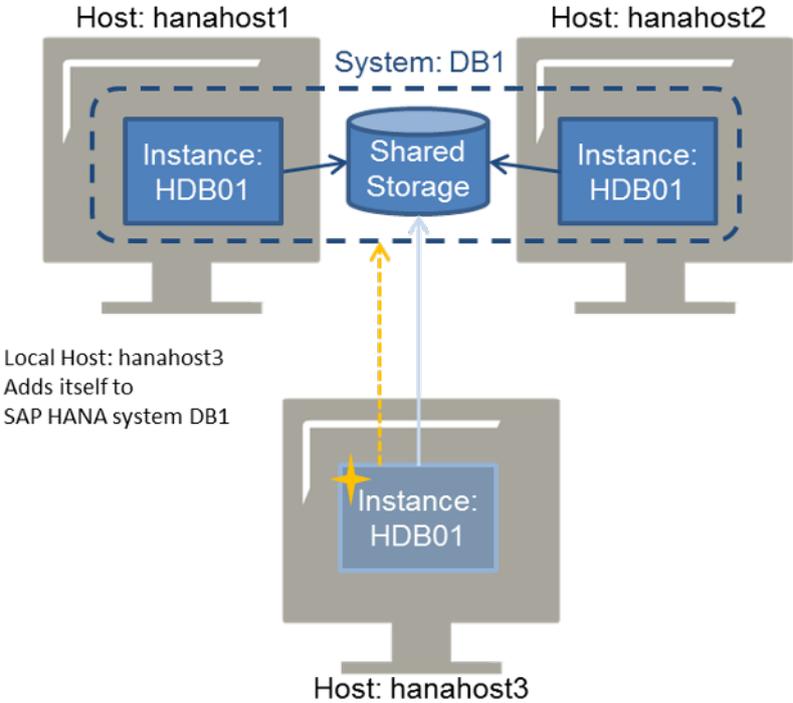


It is also possible to add multiple non-integrated hosts to the same system at one time. In the diagram below, three remote hosts (hanahost3, hanahost4, hanahost5) are added to the SAP HANA system (DB1) from a system host (hanahost1).

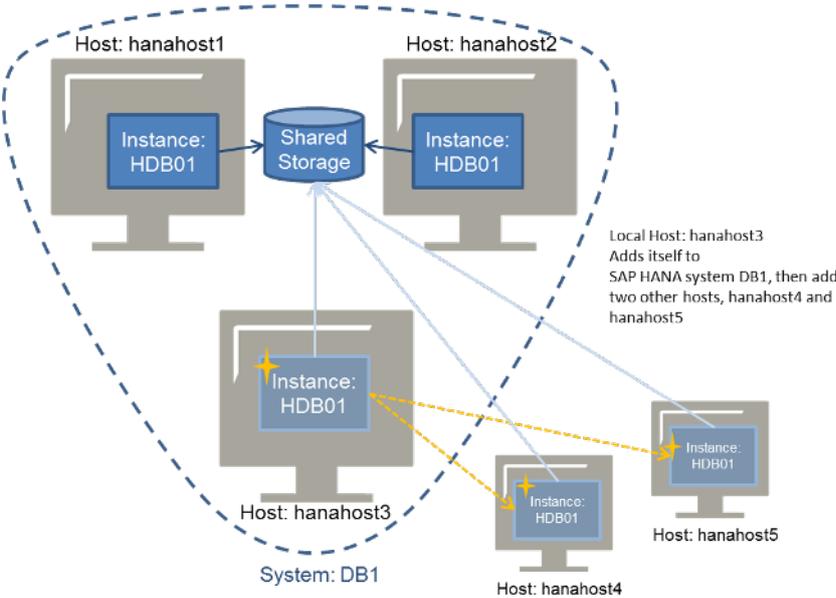


## Adding Hosts from a Non-Integrated Host

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 \[page 541\]](#)

### 5.1.2.1.3 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 \[page 532\]](#)

[Host Addition Concepts \[page 535\]](#)

[Add Hosts Using the Graphical User Interface \[page 539\]](#)

[Add Hosts Using the Command-Line Interface \[page 541\]](#)

[Add Hosts Using the Web User Interface \[page 544\]](#)

## 5.1.2.1.3.1 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.

Field Name	Description
<i>Role</i>	<p>Specifies the purpose of the SAP HANA host. SAP HANA hosts in production environments must only have one host role. However, if XS advanced runtime is installed, hosts can share multiple roles.</p> <ul style="list-style-type: none"> <li>○ Database Worker (<code>worker</code>) - A worker host (default) is used for database processing.</li> <li>○ Database Standby (<code>standby</code>) - A standby host is idle and available for fail-over in a high-availability environment.</li> <li>○ Dynamic Tiering Worker (<code>extended_storage_worker</code>) - Worker host for SAP HANA dynamic tiering</li> <li>○ Dynamic Tiering Standby (<code>extended_storage_standby</code>) - Standby host for SAP HANA dynamic tiering</li> <li>○ Accelerator for SAP ASE Worker (<code>ets_worker</code>) - Worker host for SAP HANA accelerator for SAP ASE</li> <li>○ Accelerator for SAP ASE Standby (<code>ets_standby</code>) - Standby host for SAP HANA accelerator for SAP ASE</li> <li>○ Remote Data Sync (<code>rdsync</code>) - Host for SAP HANA remote data sync</li> <li>○ Smart Data Streaming (<code>streaming</code>) - Host for SAP HANA smart data streaming</li> <li>○ XS advanced runtime worker - Host for SAP HANA XS advanced runtime</li> <li>○ XS advanced runtime standby - Standby host for SAP HANA XS advanced runtime</li> </ul>
<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><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>
<i>Internal Network Address</i>	<p>Specifies the internal subnet address in CIDR notation.</p> <p>If you define a value other than <code>local</code>, the local interfaces will always be open.</p>
<i>Certificate Host Name</i>	Specifies the hostname used for generation of self-signed SSL certificates for the SAP Host Agent.

6. 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. If your system is SAP HANA multitenant database container (multiple-container) enabled system, the new hosts have been added to the landscape information of the system database.

### 5.1.2.1.3.2 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

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 --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
<host name>	Specifies the host name of the machine.
role	Specifies the purpose of the SAP HANA host. SAP HANA hosts in production environments must only have one host role. However, if XS advanced runtime is installed, hosts can share multiple roles. <ul style="list-style-type: none"> <li>◦ <code>worker</code> - A worker host (default) is used for database processing.</li> <li>◦ <code>standby</code> - A standby host is idle and available for failover in a high-availability environment.</li> <li>◦ <code>extended_storage_worker</code> - Worker host for SAP HANA dynamic tiering</li> <li>◦ <code>extended_storage_standby</code> - Standby host for SAP HANA dynamic tiering</li> <li>◦ <code>ets_worker</code> - Worker host for SAP HANA accelerator for SAP ASE</li> <li>◦ <code>ets_standby</code> - Standby host for SAP HANA accelerator for SAP ASE</li> <li>◦ <code>streaming</code> - Host for SAP HANA smart data streaming</li> <li>◦ <code>rdsync</code> - Host for SAP HANA remote data sync</li> <li>◦ <code>xs_worker</code> - Host for SAP HANA XS advanced runtime</li> <li>◦ <code>xs_standby</code> - Standby host for SAP HANA XS advanced runtime</li> </ul>
group	Specifies the host group ID for failover scenarios. If undefined, the host group is named "default".
storage_partition	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.  <i>global</i> - Binds the processes to all interfaces. This option does not require an internal network address entry. <i>internal</i> - 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 subnet address in CIDR notation.  If you define a value other than <i>local</i> , 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.

- 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. If your system is SAP HANA multitenant database container (multiple-container) enabled system, 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* in Related Information.

### 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 \[page 535\]](#)

[Configure SAP HANA Inter-Service Communication Using the Command-Line Interface \[page 567\]](#)

[Using the SAP HANA Platform LCM Tools \[page 502\]](#)

[nstart \[page 581\]](#)

### 5.1.2.1.3.3 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 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**.
    - 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* in Related Information.

- You are logged on as the system administrator user `<sid>adm`.

You should verify that the following additional prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from the SAP HANA studio.

- The SAP HANA studio revision is 120 or higher.
- For Linux:
  - The system property `org.eclipse.swt.browser.XULRunnerPath` should be set in `hdbstudio.ini` to point to the path of XULRunner, for example:
 

```
-Dorg.eclipse.swt.browser.XULRunnerPath=<path to xulrunner>.
```

 This `hdbstudio.ini` file is located in the same folder as the executable that is used to start the SAP HANA studio. For Linux, the default location is `hana/shared/<SID>/hdbstudio..`

## Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
<b>Web Browser</b>	<p>Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser:</p> <pre>https://&lt;hostname&gt;:1129/lmsl/HDBLCM/&lt;SID&gt;/index.html</pre> <div style="background-color: #fff9c4; padding: 5px; margin-top: 10px;"> <p><b>i Note</b></p> <p>The URL is case sensitive. Make sure you enter upper and lower case letters correctly.</p> </div>
<b>SAP HANA Studio</b>	<ol style="list-style-type: none"> <li>1. Start the SAP HANA studio.</li> <li>2. In the SAP HANA studio, add the SAP HANA system.</li> <li>3. Open the context menu (right-mouse click) in the <i>Systems</i> view, and select <i>Add System</i>. For more information about adding a system, see <i>Add an SAP HANA System</i> in the <i>SAP HANA Administration Guide</i> in Related Information.</li> <li>4. In the SAP HANA studio, log on to the system.</li> <li>5. From the context menu of the selected system, select <b>Lifecycle Management</b> <b>Platform Lifecycle Management</b> <b>SAP HANA Platform Lifecycle Management</b>.</li> </ol>
<b>SAP HANA Cockpit</b>	<ol style="list-style-type: none"> <li>1. Enter the SAP HANA cockpit URL in your browser. The URL depends on whether you are connecting to a single-container system or to a database in a multiple-container system. A single-container system is accessed through the URL: <code>http://&lt;host_FQDN&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code> For more information about the URLs in multiple-container systems, see <i>Configure HTTP Access to Multitenant Database Containers</i> in the <i>SAP HANA Administration Guide</i> in Related Information.           <div style="background-color: #fff9c4; padding: 5px; margin-top: 10px;"> <p><b>i Note</b></p> <p>FQDN = fully qualified domain name</p> </div> </li> <li>2. The <i>SAP HANA Platform Lifecycle Management</i> tiles are visible on the homepage of the SAP HANA cockpit. If they are not, you can add them from the <i>SAP HANA Platform Lifecycle Management</i> tile catalog. For more information, see <i>Customizing the Homepage of SAP HANA Cockpit</i>.</li> </ol>

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
<b>import_xs_content</b>	Imports SAP HANA XS advanced runtime content.
<b>Install or Update SAP Host Agent</b>	Installs or updates SAP Host Agent.
<b>Do Not Start Added Hosts</b>	Does not start hosts after addition.
<b>Do Not Modify 'etc/sudoers' File</b>	Prevents the file <code>/etc/sudoers</code> from being modified.
<b>Timeouts</b>	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. SAP HANA hosts in production environments must only have one host role. However, if XS advanced runtime is installed, hosts can share multiple roles.</p> <ul style="list-style-type: none"> <li>○ Database Worker (worker) - A worker host (default) is used for database processing.</li> <li>○ Database Standby (standby) - A standby host is idle and available for fail-over in a high-availability environment.</li> <li>○ Dynamic Tiering Worker (extended_storage_worker) - Worker host for SAP HANA dynamic tiering</li> <li>○ Dynamic Tiering Standby (extended_storage_standby) - Standby host for SAP HANA dynamic tiering</li> <li>○ Accelerator for SAP ASE Worker (ets_worker) - Worker host for SAP HANA accelerator for SAP ASE</li> <li>○ Accelerator for SAP ASE Standby (ets_standby) - Standby host for SAP HANA accelerator for SAP ASE</li> <li>○ Remote Data Sync (rdsync) - Host for SAP HANA remote data sync</li> <li>○ Smart Data Streaming (streaming) - Host for SAP HANA smart data streaming</li> <li>○ XS advanced runtime worker - Host for SAP HANA XS advanced runtime</li> <li>○ XS advanced runtime standby - Standby host for SAP HANA XS advanced runtime</li> </ul>
<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>	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 subnet address in CIDR notation.  If you define a value other than <code>local</code> , the local interfaces will always be open.

- 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. If your system is SAP HANA multitenant database container (multiple-container) enabled system, the new hosts have been added to the landscape information of the system database.

### 5.1.2.1.4 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 \[page 548\]](#)

[Remove Hosts Using the Command-Line Interface \[page 549\]](#)

[Remove Hosts Using the Web User Interface \[page 551\]](#)

## 5.1.2.1.4.1 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. Move the host-specific content of SAP HANA system to other hosts. This procedure only applies to a single-container system. If your system is configured as a multiple-container system, you have to remove tenant-specific services first and then continue with step 2. For more information, see *Remove a Service from a Tenant Database* in the *SAP HANA Multitenant Database Containers Operations Guide*.
  - a. Start the SAP HANA studio.
  - b. Right-click the affected SAP HANA system and select **Configuration and Monitoring** **Open Administration** **Landscape** **Hosts**.
  - c. Right-click the affected host name and select *Remove Host...*  
A description is displayed, select *Yes*.
  - d. Press the refresh button in order to see the *Removal Status* 'REORG FINISHED'.
2. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblcml
```

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

3. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

```
./hdblcmgui
```

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 (<sid>adm) 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 (<sid>adm) 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 \[page 502\]](#)

[Host Addition Concepts \[page 535\]](#)

[Remove a Service from a Tenant Database \[page 138\]](#)

### 5.1.2.1.4.2 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. Move the host-specific content of SAP HANA system to other hosts. This procedure only applies to a single-container system. If your system is configured as a multiple-container system, you have to remove tenant-specific services first and then continue with step 2. For more information, see *Remove a Service from a Tenant Database* in the *SAP HANA Multitenant Database Containers Operations Guide*.
  - a. Start the SAP HANA studio.
  - b. Right-click the affected SAP HANA system and select **Configuration and Monitoring > Open Administration > Landscape > Hosts**.
  - c. Right-click the affected host name and select *Remove Host...*  
A description is displayed, select *Yes*.
  - d. Press the refresh button in order to see the *Removal Status* 'REORG FINISHED'.
2. Change to the SAP HANA resident HDBLCM directory:

```
cd <sapmnt>/<SID>/hdblcmlcm
```

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

3. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblcmlcm
```

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 (<sid>adm) 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 (<sid>adm) 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 \[page 502\]](#)

[Host Addition Concepts \[page 535\]](#)

### 5.1.2.1.4.3 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**.
  - 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* in Related Information.

- You are logged on as the system administrator user `<sid>adm`.

You should verify that the following additional prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from the SAP HANA studio.

- The SAP HANA studio revision is 120 or higher.
- For Linux:

- The system property `org.eclipse.swt.browser.XULRunnerPath` should be set in `hdbstudio.ini` to point to the path of XULRunner, for example:  
`-Dorg.eclipse.swt.browser.XULRunnerPath=<path to xulrunner>`.  
 This `hdbstudio.ini` file is located in the same folder as the executable that is used to start the SAP HANA studio. For Linux, the default location is `hana/shared/<SID>/hdbstudio..`

### 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. Move the host-specific content of SAP HANA system to other hosts. This procedure only applies to a single-container system. If your system is configured as a multiple-container system, you have to remove tenant-specific services first and then continue with step 2. For more information, see *Remove a Service from a Tenant Database* in the *SAP HANA Multitenant Database Containers Operations Guide*.
  - a. Start the SAP HANA studio.
  - b. Right-click the affected SAP HANA system and select **► Configuration and Monitoring ► Open Administration ► Landscape ► Hosts ►**.
  - c. Right-click the affected host name and select *Remove Host...*  
 A description is displayed, select *Yes*.
  - d. Press the refresh button in order to see the *Removal Status* 'REORG FINISHED'.
2. Access the SAP HANA HDBLCM Web user interface.

Option	Description
<b>Web Browser</b>	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: <code>https://&lt;hostname&gt;:1129/lmsl/HDBLCM/&lt;SID&gt;/index.html</code>  <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b> The URL is case sensitive. Make sure you enter upper and lower case letters correctly.</p> </div>
<b>SAP HANA Studio</b>	<ol style="list-style-type: none"> <li>1. Start the SAP HANA studio.</li> <li>2. In the SAP HANA studio, add the SAP HANA system.</li> <li>3. Open the context menu (right-mouse click) in the <i>Systems</i> view, and select <i>Add System</i>. For more information about adding a system, see <i>Add an SAP HANA System</i> in the <i>SAP HANA Administration Guide</i> in Related Information.</li> <li>4. In the SAP HANA studio, log on to the system.</li> <li>5. From the context menu of the selected system, select <b>► Lifecycle Management ► Platform Lifecycle Management ► SAP HANA Platform Lifecycle Management ►</b>.</li> </ol>
<b>SAP HANA Cockpit</b>	<ol style="list-style-type: none"> <li>1. Enter the SAP HANA cockpit URL in your browser. The URL depends on whether you are connecting to a single-container system or to a database in a multiple-container system.</li> </ol>

Option	Description
	<p>A single-container system is accessed through the URL: <code>http://&lt;host_FQDN&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code></p> <p>For more information about the URLs in multiple-container systems, see <i>Configure HTTP Access to Multitenant Database Containers</i> in the <i>SAP HANA Administration Guide</i> in Related Information.</p> <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>FQDN = fully qualified domain name</p> </div> <p>2. The <i>SAP HANA Platform Lifecycle Management</i> tiles are visible on the homepage of the SAP HANA cockpit. If they are not, you can add them from the <i>SAP HANA Platform Lifecycle Management</i> tile catalog. For more information, see <i>Customizing the Homepage of SAP HANA Cockpit</i>.</p>

3. Select the *Remove Hosts* tile.
4. 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
<b>Do Not Remove XS Advanced OS Users</b>	Prevents the XS advanced runtime OS Users from being removed.
<b>Do Not Modify 'etc/sudoers' File</b>	Prevents the file <code>/etc/sudoers</code> from being modified.
<b>Timeouts</b>	Sets customized timeouts ( <code>start_instance</code> , <code>start_service</code> , <code>stop_instance</code> , <code>stop_service</code> ).

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>&lt;sid&gt;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>&lt;sid&gt;adm</code> ) from being removed.

6. Enter the relevant credentials. Then select *Next*.
7. Review the summary, and select *Run* to finalize the configuration.

## Related Information

[SAPUI5 Developer Guide](#)

[Using the SAP HANA Platform LCM Tools \[page 502\]](#)

[Host Addition Concepts \[page 535\]](#)

[Remove a Service from a Tenant Database \[page 138\]](#)

---

## 5.1.2.2 Configuring Host Roles

It is possible to add and remove host roles after installation in a single-host or multiple-host SAP HANA system.

Before adding a host role, it is important to review multiple-host system concepts, and also review the SAP HANA database lifecycle manager add host technology 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 \[page 554\]](#)

[Removing Host Roles \[page 560\]](#)

## 5.1.2.2.1 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 \[page 554\]](#)

[Add Host Roles Using the Command-Line Interface \[page 556\]](#)

[Add Host Roles Using the Web User Interface \[page 557\]](#)

### 5.1.2.2.1.1 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>/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 [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
<a href="#">Role</a>	<p>Specifies the purpose of the SAP HANA host. SAP HANA hosts in production environments must only have one host role. However, if XS advanced runtime is installed, hosts can share multiple roles.</p> <ul style="list-style-type: none"> <li>○ Database Worker (worker) - A worker host (default) is used for database processing.</li> <li>○ Database Standby (standby) - A standby host is idle and available for fail-over in a high-availability environment.</li> <li>○ Dynamic Tiering Worker (extended_storage_worker) - Worker host for SAP HANA dynamic tiering</li> <li>○ Dynamic Tiering Standby (extended_storage_standby) - Standby host for SAP HANA dynamic tiering</li> <li>○ Accelerator for SAP ASE Worker (ets_worker) - Worker host for SAP HANA accelerator for SAP ASE</li> <li>○ Accelerator for SAP ASE Standby (ets_standby) - Standby host for SAP HANA accelerator for SAP ASE</li> <li>○ Remote Data Sync (rdsync) - Host for SAP HANA remote data sync</li> <li>○ Smart Data Streaming (streaming) - Host for SAP HANA smart data streaming</li> <li>○ XS advanced runtime worker - Host for SAP HANA XS advanced runtime</li> <li>○ XS advanced runtime standby - Standby host for SAP HANA XS advanced runtime</li> </ul>

5. Review the summary, and select [Run](#) to finalize the configuration.

## 5.1.2.2.1.2 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. SAP HANA hosts in production environments must only have one host role. However, if XS advanced runtime is installed, hosts can share multiple roles.</p> <ul style="list-style-type: none"><li>◦ <code>worker</code> - A worker host (default) is used for database processing.</li><li>◦ <code>standby</code> - A standby host is idle and available for failover in a high-availability environment.</li><li>◦ <code>extended_storage_worker</code> - Worker host for SAP HANA dynamic tiering</li><li>◦ <code>extended_storage_standby</code> - Standby host for SAP HANA dynamic tiering</li><li>◦ <code>ets_worker</code> - Worker host for SAP HANA accelerator for SAP ASE</li><li>◦ <code>ets_standby</code> - Standby host for SAP HANA accelerator for SAP ASE</li><li>◦ <code>streaming</code> - Host for SAP HANA smart data streaming</li><li>◦ <code>rdsync</code> - Host for SAP HANA remote data sync</li><li>◦ <code>xs_worker</code> - Host for SAP HANA XS advanced runtime</li><li>◦ <code>xs_standby</code> - Standby host for SAP HANA XS advanced runtime</li></ul>

5. Enter the required credentials.
6. Review the summary, and select **y** to finalize the configuration.

### 5.1.2.2.1.3 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**.
    - 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* in Related Information.

- You are logged on as the system administrator user `<sid>adm`.

You should verify that the following additional prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from the SAP HANA studio.

- The SAP HANA studio revision is 120 or higher.
- For Linux:
  - The system property `org.eclipse.swt.browser.XULRunnerPath` should be set in `hdbstudio.ini` to point to the path of XULRunner, for example:  
`-Dorg.eclipse.swt.browser.XULRunnerPath=<path to xulrunner>`.  
This `hdbstudio.ini` file is located in the same folder as the executable that is used to start the SAP HANA studio. For Linux, the default location is `hana/shared/<SID>/hdbstudio..`

## Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
<b>Web Browser</b>	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser:

Option	Description
	<p>https://&lt;hostname&gt;:1129/lmsl/HDBLCM/&lt;SID&gt;/index.html</p> <p><b>i Note</b> The URL is case sensitive. Make sure you enter upper and lower case letters correctly.</p>
<b>SAP HANA Studio</b>	<ol style="list-style-type: none"> <li>1. Start the SAP HANA studio.</li> <li>2. In the SAP HANA studio, add the SAP HANA system.</li> <li>3. Open the context menu (right-mouse click) in the <i>Systems</i> view, and select <i>Add System</i>. For more information about adding a system, see <i>Add an SAP HANA System</i> in the <i>SAP HANA Administration Guide</i> in Related Information.</li> <li>4. In the SAP HANA studio, log on to the system.</li> <li>5. From the context menu of the selected system, select <b>Lifecycle Management</b> <b>Platform Lifecycle Management</b> <b>SAP HANA Platform Lifecycle Management</b>.</li> </ol>
<b>SAP HANA Cockpit</b>	<ol style="list-style-type: none"> <li>1. Enter the SAP HANA cockpit URL in your browser. The URL depends on whether you are connecting to a single-container system or to a database in a multiple-container system. A single-container system is accessed through the URL: <code>http://&lt;host_FQDN&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code> For more information about the URLs in multiple-container systems, see <i>Configure HTTP Access to Multitenant Database Containers</i> in the <i>SAP HANA Administration Guide</i> in Related Information. <p><b>i Note</b> FQDN = fully qualified domain name</p> </li> <li>2. The <i>SAP HANA Platform Lifecycle Management</i> tiles are visible on the homepage of the SAP HANA cockpit. If they are not, you can add them from the <i>SAP HANA Platform Lifecycle Management</i> tile catalog. For more information, see <i>Customizing the Homepage of SAP HANA Cockpit</i>.</li> </ol>

2. Select the *Add 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
<b>Do Not Start Hosts After Addition of Roles</b>	Does not start hosts after addition of roles.
<b>Do Not Modify 'etc/sudoers' File</b>	Prevents the file <code>/etc/sudoers</code> from being modified.
<b>Timeouts</b>	Sets customized timeouts ( <code>start_instance</code> , <code>start_service</code> , <code>stop_instance</code> , <code>stop_service</code> ).

4. Select the hosts to which you would like to assign additional roles.
5. 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. SAP HANA hosts in production environments must only have one host role. However, if XS advanced runtime is installed, hosts can share multiple roles.</p> <ul style="list-style-type: none"> <li>○ <code>worker</code> - A worker host (default) is used for database processing.</li> <li>○ <code>standby</code> - A standby host is idle and available for failover in a high-availability environment.</li> <li>○ <code>extended_storage_worker</code> - Worker host for SAP HANA dynamic tiering</li> <li>○ <code>extended_storage_standby</code> - Standby host for SAP HANA dynamic tiering</li> <li>○ <code>ets_worker</code> - Worker host for SAP HANA accelerator for SAP ASE</li> <li>○ <code>ets_standby</code> - Standby host for SAP HANA accelerator for SAP ASE</li> <li>○ <code>streaming</code> - Host for SAP HANA smart data streaming</li> <li>○ <code>rdsync</code> - Host for SAP HANA remote data sync</li> <li>○ <code>xs_worker</code> - Host for SAP HANA XS advanced runtime</li> <li>○ <code>xs_standby</code> - Standby host for SAP HANA XS advanced runtime</li> </ul>

6. Enter the required credentials.
7. Review the summary, and select [Add Roles](#) to finalize the configuration.

## Related Information

[SAPUI5 Developer Guide](#)

### 5.1.2.2 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 \[page 561\]](#)

[Remove Host Roles Using the Command-Line Interface \[page 562\]](#)

[Remove Host Roles Using the Web User Interface \[page 563\]](#)

## 5.1.2.2.1 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>/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 [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.

## 5.1.2.2.2 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>/hdblc
```

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

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblc --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.

## 5.1.2.2.3 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](#).
    - 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* in Related Information.

- You are logged on as the system administrator user `<sid>adm`.

You should verify that the following additional prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from the SAP HANA studio.

- The SAP HANA studio revision is 120 or higher.
- For Linux:
  - The system property `org.eclipse.swt.browser.XULRunnerPath` should be set in `hdbstudio.ini` to point to the path of XULRunner, for example:  
`-Dorg.eclipse.swt.browser.XULRunnerPath=<path to xulrunner>`.

This `hdbstudio.ini` file is located in the same folder as the executable that is used to start the SAP HANA studio. For Linux, the default location is `hana/shared/<SID>/hdbstudio..`

## Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
<b>Web Browser</b>	<p>Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser:</p> <p><code>https://&lt;hostname&gt;:1129/lmsl/HDBLCM/&lt;SID&gt;/index.html</code></p> <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>The URL is case sensitive. Make sure you enter upper and lower case letters correctly.</p> </div>
<b>SAP HANA Studio</b>	<ol style="list-style-type: none"> <li>1. Start the SAP HANA studio.</li> <li>2. In the SAP HANA studio, add the SAP HANA system.</li> <li>3. Open the context menu (right-mouse click) in the <i>Systems</i> view, and select <i>Add System</i>. For more information about adding a system, see <i>Add an SAP HANA System</i> in the <i>SAP HANA Administration Guide</i> in Related Information.</li> <li>4. In the SAP HANA studio, log on to the system.</li> <li>5. From the context menu of the selected system, select <b>Lifecycle Management</b> <b>Platform Lifecycle Management</b> <b>SAP HANA Platform Lifecycle Management</b>.</li> </ol>
<b>SAP HANA Cockpit</b>	<ol style="list-style-type: none"> <li>1. Enter the SAP HANA cockpit URL in your browser. The URL depends on whether you are connecting to a single-container system or to a database in a multiple-container system. A single-container system is accessed through the URL: <code>http://&lt;host_FQDN&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code> For more information about the URLs in multiple-container systems, see <i>Configure HTTP Access to Multitenant Database Containers</i> in the <i>SAP HANA Administration Guide</i> in Related Information. <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>FQDN = fully qualified domain name</p> </div> </li> <li>2. The <i>SAP HANA Platform Lifecycle Management</i> tiles are visible on the homepage of the SAP HANA cockpit. If they are not, you can add them from the <i>SAP HANA Platform Lifecycle Management</i> tile catalog. For more information, see <i>Customizing the Homepage of SAP HANA Cockpit</i>.</li> </ol>

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
<b>Do Not Remove XS Advanced OS Users</b>	Prevents the XS advanced runtime OS Users from being removed.

Option	Description
<b>Do Not Start Hosts After Removal of Roles</b>	Does not start hosts after removal of roles.
<b>Do Not Modify 'etc/sudoers' File</b>	Prevents the file <code>/etc/sudoers</code> from being modified.
<b>Timeouts</b>	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](#)

### 5.1.2.3 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 \[page 566\]](#)

[Configure SAP HANA Inter-Service Communication Using the Command-Line Interface \[page 567\]](#)

[Configure SAP HANA Inter-Service Communication Using the Web User Interface \[page 569\]](#)

### 5.1.2.3.1 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>/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 *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.</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 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 subnet 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 Update Master*, *SAP HANA Security Guide*, and *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]`.

### 5.1.2.3.2 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 .
4. Define the required parameters.

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> <p><i>local</i> - Opens the communication ports for internal usage on the local interfaces. This configuration is only an option for single 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 <i>local</i>, the local interfaces will always be open.</p>
<i>Internal Network Address</i>	Specifies the internal subnet address in CIDR notation.

For more information about parameters for the *configure\_internal\_network* action, see Related Information.

5. 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* in 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 SAP HANA resident HDBLCM directory of the SAP HANA system, execute the following command:

```
./hdblcm --action=configure_internal_network --help
```

### Example

The following example configures the internal network communication with internal interface:

```
./hdblcm --action=configure_internal_network --listen_interface=internal --  
internal_address=10.66.8/21
```

## Related Information

[Configuring the Network for Multiple Hosts \[page 1000\]](#)

[nstart \[page 581\]](#)

[Using the SAP HANA Platform LCM Tools \[page 502\]](#)

### 5.1.2.3.3 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](#) .
  - 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* in Related Information.

- You are logged on as the system administrator user `<sid>adm`.

You should verify that the following additional prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from the SAP HANA studio.

- The SAP HANA studio revision is 120 or higher.
- For Linux:
  - The system property `org.eclipse.swt.browser.XULRunnerPath` should be set in `hdbstudio.ini` to point to the path of XULRunner, for example:  
`-Dorg.eclipse.swt.browser.XULRunnerPath=<path to xulrunner>`.  
 This `hdbstudio.ini` file is located in the same folder as the executable that is used to start the SAP HANA studio. For Linux, the default location is `hana/shared/<SID>/hdbstudio..`

## Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
<b>Web Browser</b>	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: <code>https://&lt;hostname&gt;:1129/lmsl/HDBLCM/&lt;SID&gt;/index.html</code>
	<b>i</b> Note The URL is case sensitive. Make sure you enter upper and lower case letters correctly.
<b>SAP HANA Studio</b>	<ol style="list-style-type: none"> <li>1. Start the SAP HANA studio.</li> <li>2. In the SAP HANA studio, add the SAP HANA system.</li> <li>3. Open the context menu (right-mouse click) in the <i>Systems</i> view, and select <i>Add System</i>. For more information about adding a system, see <i>Add an SAP HANA System</i> in the <i>SAP HANA Administration Guide</i> in Related Information.</li> <li>4. In the SAP HANA studio, log on to the system.</li> <li>5. From the context menu of the selected system, select <b>Lifecycle Management</b> <b>Platform Lifecycle Management</b> <b>SAP HANA Platform Lifecycle Management</b>.</li> </ol>
<b>SAP HANA Cockpit</b>	<ol style="list-style-type: none"> <li>1. Enter the SAP HANA cockpit URL in your browser. The URL depends on whether you are connecting to a single-container system or to a database in a multiple-container system. A single-container system is accessed through the URL: <code>http://&lt;host_FQDN&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code></li> </ol>

Option	Description
	<p>For more information about the URLs in multiple-container systems, see <i>Configure HTTP Access to Multitenant Database Containers</i> in the <i>SAP HANA Administration Guide</i> in Related Information.</p> <p><b>i Note</b></p> <p>FQDN = fully qualified domain name</p> <p>2. The <i>SAP HANA Platform Lifecycle Management</i> tiles are visible on the homepage of the SAP HANA cockpit. If they are not, you can add them from the <i>SAP HANA Platform Lifecycle Management</i> tile catalog. For more information, see <i>Customizing the Homepage of SAP HANA Cockpit</i>.</p>

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
<b>nostart</b>	Prevents the SAP HANA system from being started.
<b>Timeouts</b>	Sets customized timeouts ( <code>start_instance</code> , <code>start_service</code> , <code>stop_instance</code> , <code>stop_service</code> ).

4. Provide the password of the `<sid>adm` user, then select [Next](#).
5. Specify values for the following fields:

Field Name	Description
<a href="#">Inter-Service Communication</a>	<p>Specifies the listen interface for the internal network communication.</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 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>
<a href="#">Internal Network Address</a>	Specifies the internal subnet 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 Update Master*, *SAP HANA Security Guide*, and *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 \[page 70\]](#)

### 5.1.2.4 Converting an SAP HANA System to Support Multitenant Database Containers

You can convert an SAP HANA system to support multitenant database containers using the SAP HANA database lifecycle manager (HDBLCM) resident program. Converting an SAP HANA system to a multiple-container system is permanent and cannot be reversed.

If your system was installed in single-container mode, you can still implement multitenant database containers by converting the system to a multiple-container 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, and connection properties (port configuration). However, it does **not** contain the backup history and the system license.

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 Multitenant Database Containers Using the Graphical User Interface \[page 573\]](#)

[Convert to Multitenant Database Containers Using the Command-Line Interface \[page 575\]](#)

[Convert to Multitenant Database Containers Using the Web User Interface \[page 577\]](#)

[Perform an Offline Conversion \[page 582\]](#)

## 5.1.2.4.1 Convert to Multitenant Database Containers Using the Graphical User Interface

You can convert an SAP HANA system to support multitenant database containers using the SAP HANA database lifecycle manager (HDBLCM) resident program in the graphical user interface. Converting an SAP HANA system to a multiple-container 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 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 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 *Convert to Multitenant Database Containers* 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 multiple-container 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 studio 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 and the original license).
- Its trace and configuration files are now stored at the following location: `/usr/sap/<SID>/HDB<instance>/host/trace/DB_<database_name>`.

### 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*.

### 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 Multiple-Container Systems* in the *SAP HANA Administration Guide*.

- Update the configuration of the Web Dispatcher. For more information, see *Configure HTTP Access to Multitenant Database Containers* in the *SAP HANA Administration Guide*.

## Related Information

[SAP Note 1917938](#)

[Password Policy Configuration Options \[page 655\]](#)

[Create a Tenant Database \[page 111\]](#)

[Deploy a Delivery Unit Archive \(\\*.tgz\) \[page 632\]](#)

[Install a Permanent License \[page 224\]](#)

[Creating Backups \[page 920\]](#)

[Configuration Parameters in Multiple-Container Systems \[page 213\]](#)

## 5.1.2.4.2 Convert to Multitenant Database Containers Using the Command-Line Interface

You can convert an SAP HANA system to support multitenant database containers using the SAP HANA database lifecycle manager (HDBLCM) resident program in the command-line interface. Converting an SAP HANA system to a multiple-container 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 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 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 --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 multiple-container 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 studio 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 and the original license).
- Its trace and configuration files are now stored at the following location: `/usr/sap/<SID>/HDB<instance>/host/trace/DB_<database_name>`.

### 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*.

### 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 Multiple-Container Systems* in the *SAP HANA Administration Guide*.

- Update the configuration of the Web Dispatcher. For more information, see *Configure HTTP Access to Multitenant Database Containers* in the *SAP HANA Administration Guide*.

## Related Information

[SAP Note 1917938](#)

[Password Policy Configuration Options \[page 655\]](#)

[Create a Tenant Database \[page 111\]](#)

[Deploy a Delivery Unit Archive \(\\*.tgz\) \[page 632\]](#)

[Install a Permanent License \[page 224\]](#)

[Creating Backups \[page 920\]](#)

[Configuration Parameters in Multiple-Container Systems \[page 213\]](#)

---

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

[import\\_content \[page 581\]](#)

[nostart \[page 581\]](#)

[nostart\\_tenant\\_db \[page 581\]](#)

### 5.1.2.4.3 Convert to Multitenant Database Containers Using the Web User Interface

You can convert an SAP HANA system to support multitenant database containers using the SAP HANA database lifecycle manager Web user interface. Converting an SAP HANA system to a multiple-container 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 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**.
    - 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* in Related Information.

- You are logged on as the system administrator user `<sid>adm`.

You should verify that the following additional prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from the SAP HANA studio.

- The SAP HANA studio revision is 120 or higher.
- For Linux:
  - The system property `org.eclipse.swt.browser.XULRunnerPath` should be set in `hdbstudio.ini` to point to the path of XULRunner, for example:  
`-Dorg.eclipse.swt.browser.XULRunnerPath=<path to xulrunner>`.  
This `hdbstudio.ini` file is located in the same folder as the executable that is used to start the SAP HANA studio. For Linux, the default location is `hana/shared/<SID>/hdbstudio..`

## Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
<b>Web Browser</b>	<p>Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: <code>https://&lt;hostname&gt;:1129/lmsl/HDBLCM/&lt;SID&gt;/index.html</code></p> <p><b>i</b> Note The URL is case sensitive. Make sure you enter upper and lower case letters correctly.</p>
<b>SAP HANA Studio</b>	<ol style="list-style-type: none"><li>1. Start the SAP HANA studio.</li><li>2. In the SAP HANA studio, add the SAP HANA system.</li><li>3. Open the context menu (right-mouse click) in the <i>Systems</i> view, and select <i>Add System</i>.</li><li>4. In the SAP HANA studio, log on to the system.</li><li>5. From the context menu of the selected system, select <b>Lifecycle Management</b> <b>Platform Lifecycle Management</b> <b>SAP HANA Platform Lifecycle Management</b>.</li></ol>
<b>SAP HANA Cockpit</b>	<ol style="list-style-type: none"><li>1. Enter the SAP HANA cockpit URL in your browser.<p><b>i</b> Note FQDN = fully qualified domain name</p></li><li>2. The <i>SAP HANA Platform Lifecycle Management</i> tiles are visible on the homepage of the SAP HANA cockpit. If they are not, you can add them from the <i>SAP HANA Platform Lifecycle Management</i> tile catalog. For more information, see <i>Customizing the Homepage of SAP HANA Cockpit</i>.</li></ol>

2. Select the *Convert to Multitenant Database Containers* 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
<b>Import Delivery Units In The System Database</b>	Import Delivery Units In The System Database
<b>Do Not Start Instance After Reconfiguration</b>	Do Not Start Instance After Reconfiguration
<b>Do Not Start Tenant Database After Reconfiguration</b>	Do Not Start Tenant Database After Reconfiguration
<b>Timeouts</b>	Sets customized timeouts ( <code>start_instance</code> , <code>stop_instance</code> ).

- Provide the password of the `<sid>adm` user and `SYSTEM` user of `SYSTEMDB` user, then select *Next*.
- Review the summary, and select *Run* to finalize the configuration.

## Results

Your SAP HANA system is a multiple-container 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 studio 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 and the original license).
- Its trace and configuration 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 Multiple-Container Systems* in the *SAP HANA Administration Guide*.

- Update the configuration of the Web Dispatcher. For more information, see *Configure HTTP Access to Multitenant Database Containers* in the *SAP HANA Administration Guide*.

## Related Information

[SAPUI5 Developer Guide](#)

[SAP Note 1917938](#)

[Password Policy Configuration Options \[page 655\]](#)

[Create a Tenant Database \[page 111\]](#)

[Deploy a Delivery Unit Archive \(\\*.tgz\) \[page 632\]](#)

[Install a Permanent License \[page 224\]](#)

[Creating Backups \[page 920\]](#)

[Configuration Parameters in Multiple-Container Systems \[page 213\]](#)

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

### 5.1.2.4.4 Parameter Reference: Converting an SAP HANA System to Support Multitenant Database Containers

Parameters can be specified when converting an SAP HANA system to multitenant database containers in order to customize the configuration task.

The SAP HANA database lifecycle manager convert to multiddb 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 multiddb task with the following command:

```
./hdblcm --action=convert_to_multiddb --help
```

#### 5.1.2.4.4.1 `import_content`

Imports delivery units.

##### Syntax

In the command line, the following syntax is used:

```
--import_content [=off]
```

##### Remarks

The default for this parameter is `--import_content`.

#### 5.1.2.4.4.2 `nostart`

Prevents the SAP HANA system from being started.

##### Syntax

In the command line, the following syntax is used:

```
--nostart
```

#### 5.1.2.4.4.3 `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
```

---

## 5.1.2.4.5 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>/hdblcml/hdblcml --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](#) 

## 5.1.3 Changing the SAP HANA System

After installation the SAP HANA system can be configured for compatibility with other SAP products or reconfigured from the original installation settings.

### Related Information

[Changing System Identifiers \[page 583\]](#)

[Reconfiguring the SAP HANA System \[page 591\]](#)

[Managing SAP HANA System Components \[page 596\]](#)

[Check the Installation Using the Command-Line Interface \[page 597\]](#)

### 5.1.3.1 Changing System Identifiers

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.

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	<a href="#">Rename the SAP HANA System</a> > <a href="#">Define Host Properties</a> > <a href="#">Edit Host</a> >	<code>--action=rename_system -- hostmap=&lt;old host&gt;=&lt;new host&gt;</code>
Change the SID of an SAP HANA System	<a href="#">Rename the SAP HANA System</a> > <a href="#">Define System Properties</a> > <a href="#">Target System ID</a> >	<code>--action=rename_system -- target_sid=&lt;new sid&gt;</code>
Change the Instance Number of an SAP HANA System	<a href="#">Rename the SAP HANA System</a> > <a href="#">Define System Properties</a> > <a href="#">Target Instance Number</a> >	<code>--action=rename_system -- number=&lt;new instance number&gt;</code>

## Mounted SID Preparation

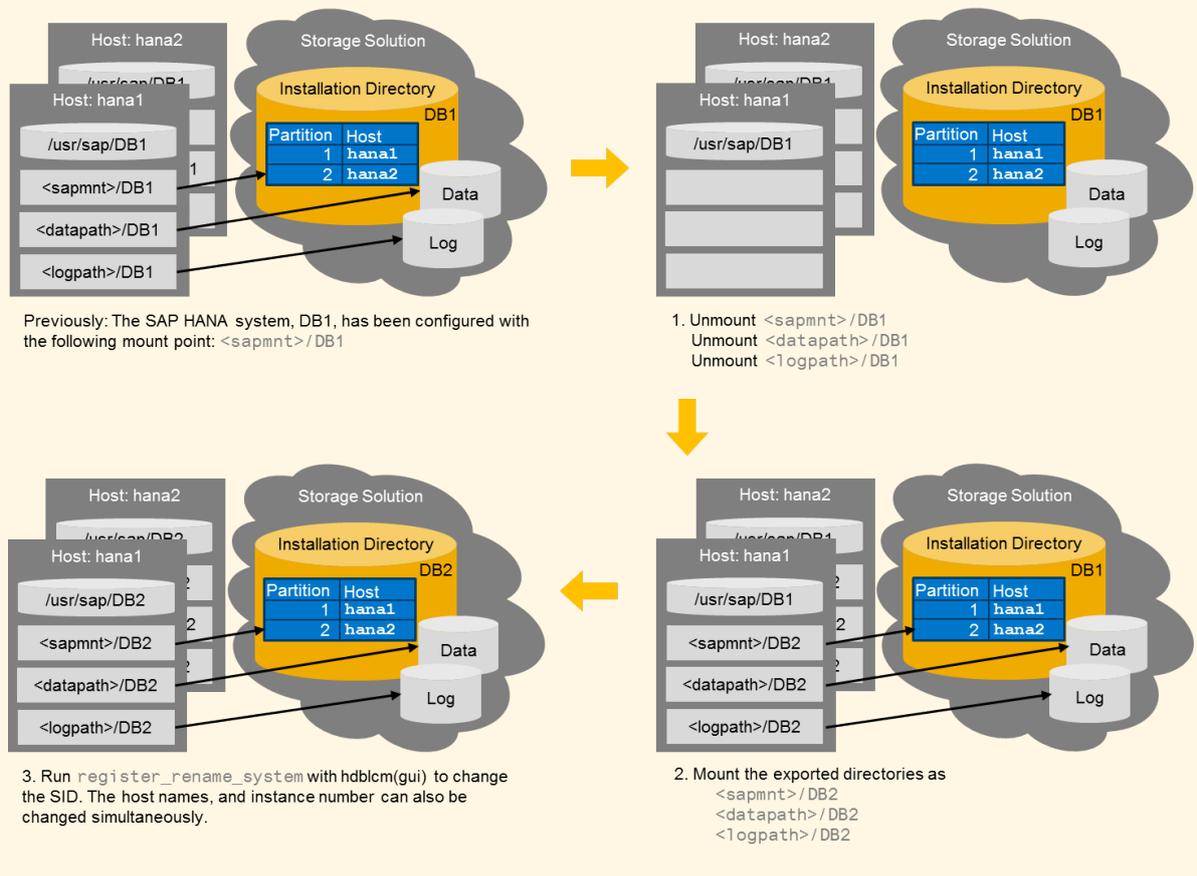
If the SID is included in the mount point, and you want to change the SID, you have to take extra preparation steps.

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



## 5.1.3.1.1 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.

### 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
```

`rename_system` action, or

- Start the tool with the `rename_system` action specified: and enter the index of the

```
./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:

```
./hdblcmgui
```

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 *and enter the index of 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:

```
./hdblcmm --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. 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*.

## Related Information

[Managing SAP HANA Licenses \[page 222\]](#)

## 5.1.3.1.2 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.

### 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 installed in multiple-container mode has one SID for the system database and all tenants. Renaming an MDC system changes the SID for the system database and all tenants. It is not possible to change the SID of individual 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:

- To change the SID using the SAP HANA database lifecycle manager command-line interface:
  - Start the command-line tool interactively:

```
./hdblcmlcm
```

and enter the index of the `rename_system` action, or

- Start the tool with the `rename_system` action specified:

```
./hdblcmlcm --action=rename_system --source_sid=<current SID> --  
target_sid=<new SID>
```

- To change the SID using the SAP HANA database lifecycle manager graphical user interface:
  1. Start the graphical user interface tool:

```
./hdblcmlcmgui
```

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.

### **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.

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:

```
./hdbclm --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. 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*.

## 5.1.3.1.3 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>/hdbclm
```

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:

```
./hdbclm
```

and enter the index of the `rename_system` action, or

- Start the tool with the `rename_system` action specified:

```
./hdbclm --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:

```
./hdbclmgui
```

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. 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*.

## Related Information

[Managing SAP HANA Licenses \[page 222\]](#)

### 5.1.3.2 Reconfiguring the SAP HANA System

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).

#### 5.1.3.2.1 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:

```
./hdblcm
```

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>/hdblcm
```

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:

```
./hdblcm
```

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*.

### **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=unregister_instance --help
```

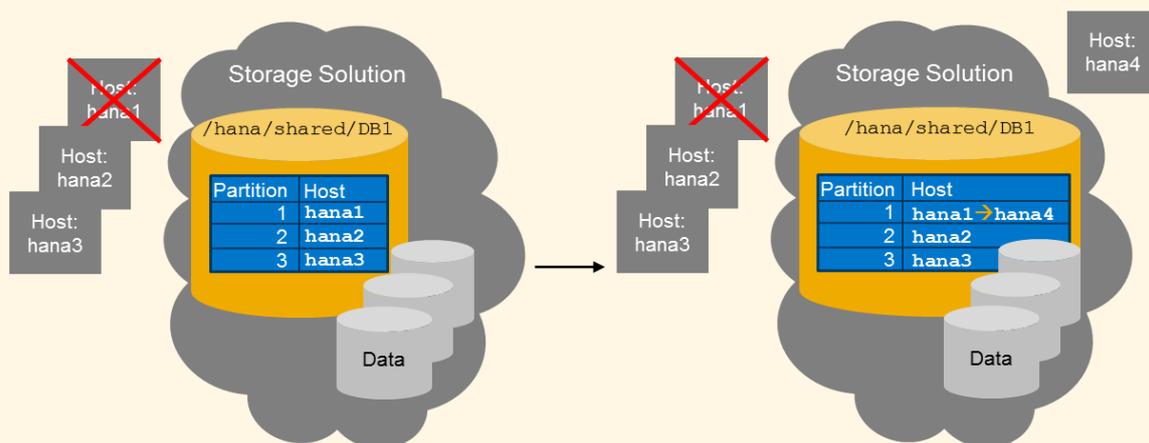
```
./hdblcm --action=unregister_system --help
```

```
./hdblcm --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.

## Related Information

[nostart \[page 581\]](#)

### 5.1.3.2.2 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.

#### 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

A system copy overwrites all users and roles 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>/hdblcmlcm
```

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:

```
./hdblcmlcm
```

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:

```
./hdblcmlcmgui
```

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](#).

### **i** Note

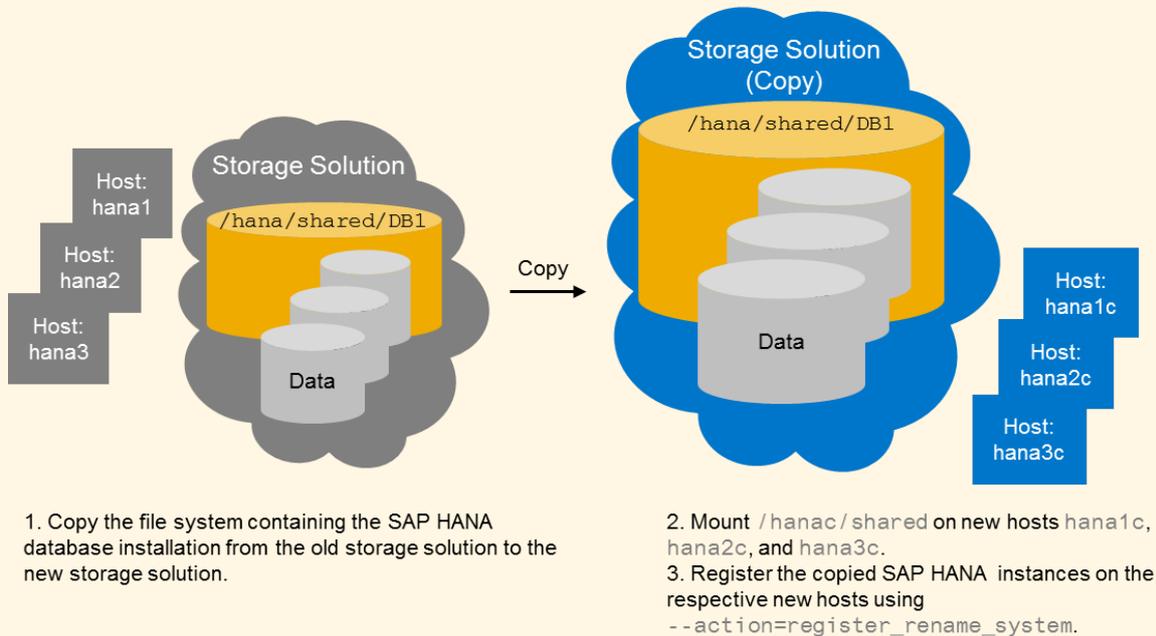
When using the command line, the options can be set interactively during configuration only if they are marked as interactive in the help description. All other options have to be specified in the command line. To call the help, in the SAP HANA resident HDBLCM directory of the SAP HANA system, execute the following command:

```
./hdblcmlcm --action=register_rename_system --help
```

## Example

The following is an example of an SAP HANA being cloned:

### SAP HANA System Clone



## 5.1.4 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.

### 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 smart data streaming
  - 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](#)

[SAP HANA Options in SAP Help Portal](#)

[Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities \[page 1360\]](#)

## 5.1.5 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.
- 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>/hdblcmm
```

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

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblcmm --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:

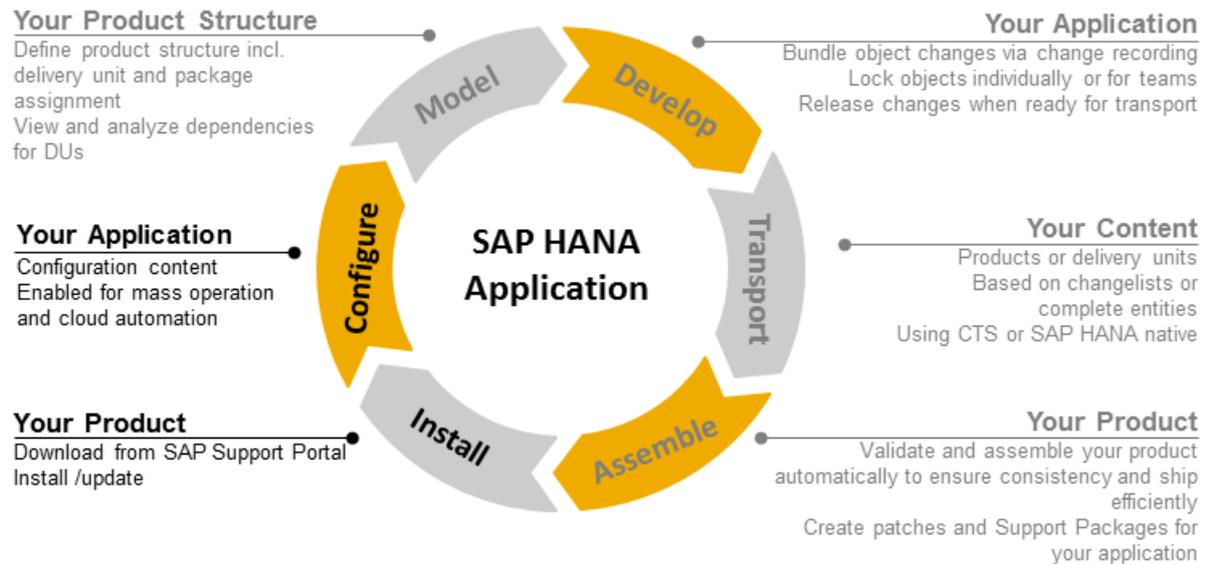
```
<sapmnt>/<SID>/global/hdb/install/support/hdbcheck.xml
```

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>_hdblcmm_check_installation_<time stamp>/hdblcmm.log
```

## 5.2 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 \[page 600\]](#)

### 5.2.1 Installing and Updating SAP HANA Products and Software Components

SAP HANA application lifecycle management provides functions for installing and updating SAP HANA products 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.
  - Start SAP HANA cockpit at the following URL: `http://<host_FQDN>:<port>/sap/hana/admin/cockpit`. Afterwards, choose the **Install Products and Software Components** tile in the **SAP HANA Application Installation and Update** group.  
For more information about how to start SAP HANA cockpit, see *Open SAP HANA Cockpit* in this *SAP HANA Administration Guide*. The link is in the *Related Information* section.

For more information about using SAP HANA Application Lifecycle Management to install and update SAP HANA products and software components, see *Installing and Updating SAP HANA Products* and *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 *Using hdbalm* and *hdbalm install Command* in the *SAP HANA Application Lifecycle Management Guide*.

## Related Information

[Installing and Updating SAP HANA Products \[page 601\]](#)

[Installing and Updating SAP HANA Software Components \[page 603\]](#)

[Installation and Update Options \[page 605\]](#)

[SAP HANA Content \[page 630\]](#)

[Open SAP HANA Cockpit \[page 23\]](#)

### 5.2.1.1 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:

- o *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.

**i** 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 \[page 605\]](#)

[Installing and Updating SAP HANA Products and Software Components \[page 600\]](#)

### 5.2.1.2 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:

- o *Status*

The following status values exist:

- o *New*

The software component is not yet installed and can be installed.

- o *Update*

The software component is already installed and can be updated to a higher version.

- o *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 `hdbalm` 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.

- o *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.

- o *Version* that is already *installed* in the system

- o *New version* that is to be installed

- o Whether the software component is part of a product that is already installed.

- o *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 \[page 605\]](#)

[Installing and Updating SAP HANA Products and Software Components \[page 600\]](#)

[SAP HANA Content \[page 630\]](#)

### 5.2.1.3 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	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: <ul style="list-style-type: none"><li>• If a previous installation operation failed, for example, because of activation errors.</li><li>• If you run continuous integration scenarios in which the same version of a software component is installed regularly.</li></ul>
Keep newer version of software component	ALLOW_KEEP_DU_NEWER_VERSION	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.

Option	Corresponding Installation Option in hdbalm	Description
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>
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 \[page 600\]](#)

## 5.2.2 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.

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)*. The link is in the *Related Information* section.

3. 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 XSA 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 log 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 <log 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 <log ID> --scv
```

## Related Information

[Prerequisites and Authorizations \[page 609\]](#)

[Installation and Update Options in XS Advanced Model \[page 614\]](#)

[Examples: Installing and Updating Products and Software Components in XS Advanced Model \[page 616\]](#)

### 5.2.2.1 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*.
- Optional when using the command line interface (CLI): The XS 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

[Setting Up Security Artifacts \[page 1154\]](#)

### 5.2.2.2 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 XSA 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 *12.2 XS CLI: Application Management* in the *SAP HANA Developer Guide For SAP HANA XS Advanced Model*. The link to the guide is in the *Related Information* section.

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 *12.2 XS CLI: Application Management* in the *SAP HANA Developer Guide For SAP HANA XS Advanced Model*. The link to the guide is in the *Related Information* section.

## Related Information

[Installing and Updating Products and Software Components in SAP HANA XS Advanced Model \[page 607\]](#)

### 5.2.2.3 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.

### 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.

### 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 \[page 607\]](#)

[Installation and Update Options in XS Advanced Model \[page 614\]](#)

[Display installed Products and Software Components in XS Advanced Model \[page 620\]](#)

## 5.2.2.4 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>'. The installation is performed only if the given archive is a product archive. Otherwise, the installation will fail.
-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> . The link is in the <i>Related Information</i> section.

Option	Description
-t <TIMEOUT>	Specify the maximum amount of time (in seconds) that the installation service must wait for the installation operation to complete
-o <VERSION_OPTION_1>[,<VERSION_OPTION_2>]	<p>Specify options which can be used to <b>override</b> the default behavior of the <code>install</code> command. The following options are available:</p> <ul style="list-style-type: none"> <li>• <b>ALLOW_PV_DOWNGRADE</b> 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.</li> <li>• <b>ALLOW_KEEP_SC_NEWER_VERSION</b> 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.</li> <li>• <b>ALLOW_SC_DOWNGRADE</b> 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.</li> </ul> <div data-bbox="718 1473 1394 1585" style="background-color: #fff9c4; padding: 5px;"> <p><b>⚠ Caution</b> Use this option carefully.</p> </div> <ul style="list-style-type: none"> <li>• <b>ALLOW_SC_SAME_VERSION</b> 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.</li> </ul>
-i, --INSTANCES <INSTANCE_1>[,<INSTANCE_2>]	By default all instances are installed; a comma-separated list of instances can be specified to limit the number of instances installed
--delete-services	Recreate changed services and/or delete discontinued services

Option	Description
<code>--delete-service-brokers</code>	Delete discontinued service brokers
<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: Plugin Management* topic in the *SAP HANA Developer Guide for SAP HANA XS Advanced Model*. The link is in the *Related Information* section.

## Related Information

[Installing and Updating Products and Software Components in SAP HANA XS Advanced Model \[page 607\]](#)  
[Examples: Installing and Updating Products and Software Components in XS Advanced Model \[page 616\]](#)

### 5.2.2.5 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                  -                    -          SUCCESS
                  SCV_A                1.0.0     SUCCESS
                  SCV_B                1.0.0     SUCCESS
-----
```

- **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                  -                    -          SUCCESS
                  SCV_A                1.5.0     SUCCESS
                  SCV_B                1.5.0     SUCCESS
-----
```

- **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                  -                -          SUCCESS
                  SCV_A            1.3.0     SUCCESS
                  SCV_B            1.3.0     SUCCESS
-----
```

- **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. To allow the installation of a higher product version, you must use the `ALLOW_PV_VERSION_UPDATE` option. In addition, to allow the installation of higher software component versions, you must use the `ALLOW_SC_VERSION_UPDATE` option.

```
xs install XSASAMPLEPRODUCT_2.0.1.zip -o ALLOW_PV_VERSION_UPDATE,
ALLOW_SC_VERSION_UPDATE
```

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                  -                -          SUCCESS
                  SCV_A            2.1.0     SUCCESS
                  SCV_B            2.1.0     SUCCESS
-----
```

- **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                  -                    -        SUCCESS
                  SCV_A                1.3.5   SUCCESS
                  SCV_B                1.3.5   SUCCESS
-----
```

- **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
```

## 5.2.2.6 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: Plugin Management* topic in *SAP HANA Developer Guide (for SAP HANA XS Advanced Model)*. The link is in the *Related Information* section.

### 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:

- a. 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.

- b. To display all software components that are installed in the current organization and space, use the `xs list-components` command.

```
xs list-components
```

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>`.

- c. 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>`. The following is an example:

```
xs list-products XSASAMPLEPRODUCT sap.com
```

### 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:

- **SUCCESS**: The software component or product instance is successfully installed.
- **BROKEN**: The software component or product instance is installed in a broken state. This status can occur, for example, if there was an error in the deploy step during installation. Product instances can get this status, for example, after one software component of the product instance was uninstalled.
- **RUNNING**: The installation of this software component or product instance is currently running.

### Example

#### Examples:

- 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     SUCCESS
            JAVA_HELLO_XSA_B    1.0.0   SUCCESS
            JAVA_HELLO_XSA_A    1.0.0   SUCCESS
3            -                  1.0     SUCCESS
            JAVA_HELLO_XSA_C    1.1.0   SUCCESS
            JAVA_HELLO_XSA_D    1.1.0   SUCCESS
-----
```

## Related Information

[Prerequisites and Authorizations \[page 609\]](#)

### 5.2.2.7 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; <b>optional</b> : 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 <code>-pv</code> option. In this case, the uninstallation is only performed if you specify a product name as <code>&lt;NAME&gt;</code> . 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 <code>-scv</code> option. In this case, the uninstallation is only performed if you specify a software component name as <code>&lt;NAME&gt;</code> . If you specify a product name, the uninstallation will fail.

Uninstallation Option	Description
<code>--ignore-scv-reuse</code>	<p>Remove the specified software component even if it is used in other installed products.</p> <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>You can use this option for uninstalling software components only.</p> </div> <p>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 <code>--ignore-scv-reuse</code> option.</p>
<code>-f</code>	Remove the specified product or software component without any system prompts or confirmation
<code>-i, --INSTANCES</code> <code>&lt;INSTANCE_1&gt;[, &lt;INSTANCE_2&gt;]</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 \[page 609\]](#)

## 5.2.3 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](http://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 \[page 626\]](#)

[Process Engine Roles \[page 626\]](#)

[Troubleshooting \[page 629\]](#)

## 5.2.3.1 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

Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. 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.

## 5.2.3.2 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 \[page 626\]](#)

[Troubleshooting \[page 629\]](#)

### 5.2.3.3 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.

## 5.3 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), can also be grouped in a DU.

For more information about SAP HANA content, see *Components Delivered as SAP HANA Content* in the *SAP HANA Security Guide*.

### 5.3.1 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 \(\\*.ZIP\) \[page 631\]](#)

[Deploy a Delivery Unit Archive \(\\*.tgz\) \[page 632\]](#)

## 5.3.2 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 \[page 600\]](#)

---

### 5.3.3 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 \[page 600\]](#)

## 6 Security Administration

Security administration represents a category of administration usually handled separately from general system administration tasks.

### [Monitoring Critical Security Settings in SAP HANA Cockpit \[page 633\]](#)

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.

### [Managing SAP HANA Users \[page 637\]](#)

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.

### [Auditing Activity in SAP HANA Systems \[page 718\]](#)

Auditing provides you with visibility on who did what in the SAP HANA database (or tried to do what) and when.

### [Managing Encryption Keys \[page 734\]](#)

SAP HANA generates unique encryption keys on installation for all mechanisms used in SAP HANA to encrypt data. However, if you received SAP HANA pre-installed from a hardware or hosting partner, you might want to change encryption keys to ensure they are not known outside your organization.

### [Managing Encryption of Data Volumes in the SAP HANA Database \[page 749\]](#)

To protect data saved to disk from unauthorized access at operating system level, the SAP HANA database supports data encryption in the persistence layer.

### [Managing Client Certificates in the SAP HANA Database \[page 758\]](#)

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.

### 6.1 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 the SAP HANA cockpit to monitor critical security settings, please refer to the *Security Configuration Checklist* in the *SAP HANA Security Guide*. This checklist provides more detailed information as well as recommendations for many settings.

---

[View Status of Security Settings \[page 634\]](#)

You can view the status of critical security settings in the SAP HANA cockpit on the tiles of the *SAP HANA Security Overview* group.

[Network Security Details \[page 635\]](#)

The *Network Security* app of the SAP HANA cockpit allows you to view important configuration settings related to secure internal SAP HANA communication and secure external SQL client communication.

## 6.1.1 View Status of Security Settings

You can view the status of critical security settings in the SAP HANA cockpit on the tiles of the *SAP HANA Security Overview* group.

### Prerequisites

- You have the privileges granted by the role `sap.hana.security.cockpit.roles::DisplaySecurityDashboard`.  
You can grant roles using the *Assign Roles* app of the SAP HANA cockpit. For more information, see *Assign Roles to a User*.
- The tiles of the *SAP HANA Security Overview* catalog are visible on the homepage of the SAP HANA cockpit. If they're not, you can add them from the tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

### Procedure

Review the security status displayed on the various tiles, drilling down to the supporting apps for more detailed information and functions.

For information about the individual tiles, see *Tile Catalog: SAP HANA Security Overview*.

### Related Information

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Assign Roles to a User \[page 717\]](#)

[Tile Catalog: SAP HANA Security Overview \[page 46\]](#)

## 6.1.2 Network Security Details

The *Network Security* app of the SAP HANA cockpit allows you to view important configuration settings related to secure internal SAP HANA communication and secure external SQL client communication.

### General Settings

Field	Description
<i>Cryptographic Provider</i>	The cryptographic service provider being used by the SAP HANA server
<i>Maximum TLS/SSL Protocol Version Accepted</i>	The maximum TLS/SSL protocol version accepted
<i>Minimum TLS/SSL Protocol Version Accepted</i>	The minimum TLS/SSL protocol version accepted
<i>Allowed TSL/SSL Cipher Suites</i>	The encryption algorithms allowed for TSL/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
<i>TSL/SSL Secured</i>	Indicates whether or not internal communication channels are secured using TSL/SSL  The following values are possible: <ul style="list-style-type: none"><li>• <i>False</i> (default)</li><li>• <i>System PKI</i></li><li>• <i>Manual configuration</i></li></ul>

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"> <li>• <i>Local network only</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.</li> <li>• <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.</li> </ul> <div style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;"> <p><b>⚠ Caution</b></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"> <li>• <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.</li> </ul>
<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

#### External JDBC/ODBC Communication

Field	Description
<i>Enforce TSL/SSL for SQL Connections</i>	<p>Indicates whether all clients communicating with the SAP HANA database via the SQL interface are required to use a secured connection</p> <p>The database refuses SQL connection attempts that don't use TSL/SSL.</p>

Field	Description
<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

## Related Information

[Configuring SAP HANA Inter-Service Communication \[page 565\]](#)

[Configuring the Network for Multiple Hosts \[page 1000\]](#)

## 6.2 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.

### **i** 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 are available for managing application users and roles. For more information, see *Maintaining the SAP HANA XS Advanced Model Run Time*.

#### [Database Users \[page 638\]](#)

Every user who wants to work with the SAP HANA database must have a database user.

#### [Operating System User <sid>adm \[page 647\]](#)

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 \[page 648\]](#)

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 \[page 669\]](#)

After successful logon, the user's authorization to perform the requested operations on the requested objects is verified.

[Provisioning Users \[page 702\]](#)

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 \[page 1139\]](#)

## 6.2.1 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.

### Standard Users

Standard users are 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, 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 (and cannot be 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 in the *Restricted User* editor of the SAP HANA studio.  
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 (`ALTER USER <username> GRANT | REVOKE ROLE PUBLIC`)
- Granting (or revoking) authorization to create objects in the user's own schema (`ALTER USER <username> GRANT | REVOKE CREATE ANY ON OWN SCHEMA`)
- Enabling (or disabling) full SQL (`ALTER USER <user_name> ENABLE CLIENT CONNECT` or enabling the corresponding option for the user in the SAP HANA cockpit)

### **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.

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

## Related Information

[Deactivate the SYSTEM User \[page 640\]](#)

### 6.2.1.1 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 addition, to ensure that the administration tool SAP HANA cockpit can be used immediately after database creation, `SYSTEM` is automatically granted several roles the first time the cockpit is opened with this user. For more information, see *Roles Granted to Database User*

---

`SYSTEM`. Note however that `SYSTEM` does not automatically have access to objects created in the SAP HANA repository.

In a system with multitenant database containers, the `SYSTEM` user of the system database 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 systems. 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`.

If the `SYSTEM` user's password is lost, you can reset it using the operating system user (<sid>adm user).

## Related Information

[Roles Granted to Database User SYSTEM \[page 56\]](#)

### 6.2.1.2 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

`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 addition, to ensure that the administration tool SAP HANA cockpit can be used immediately after database creation, `SYSTEM` is automatically granted several roles the first time the cockpit is opened with this user. For more information, see *Roles Granted to Database User SYSTEM*. Note however that `SYSTEM` does not automatically have access to objects created in the SAP HANA repository.

In a system with multitenant database containers, the `SYSTEM` user of the system database 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 systems. 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`.

## Procedure

Execute the following statement, for example, in the SQL console of the SAP HANA studio:

```
ALTER USER SYSTEM DEACTIVATE USER NOW
```

## Results

The SYSTEM user is deactivated and can no longer connect to the SAP HANA database.

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.

## 6.2.1.3 Reset the SYSTEM User's Password

If the SYSTEM user's password is lost, you can reset it as the operating system administrator by starting the index server in emergency mode.

## Prerequisites

- You cannot log on to the database as the SYSTEM because the password has been irretrievably lost.

### **i** 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

1. Log on to the server on which the master index server 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 by executing the following commands:

- o /usr/sap/<SID>/HDB<instance>/hdbenv.sh
- o /usr/sap/<SID>/HDB<instance>/exe/hdbnameserver

5. In a new session, start the compile server by executing the following commands:

- o /usr/sap/<SID>/HDB<instance>/hdbenv.sh
- o /usr/sap/<SID>/HDB<instance>/exe/hdbcompileserver

6. In a new session, start the index server by executing the following commands:

- o /usr/sap/<SID>/HDB<instance>/hdbenv.sh
- o /usr/sap/<SID>/HDB<instance>/exe/hdbindexserver -resetUserSystem

### **i** Note

In a scale-out system, you only need to execute the commands on the master index server.

After some start-up notifications, the prompt resetting of user SYSTEM - new password appears:

```
login as: [redacted]  
Using keyboard-interactive authentication.  
Password:  
Last login: [redacted] from [redacted]  
[redacted]:/usr/sap/[redacted] > /usr/sap/[redacted]/hdbenv.sh  
[redacted]:/usr/sap/[redacted] > /usr/sap/[redacted]/exe/hdbindexserver -resetUserSystem  
Starting interactive mode for resetting user SYSTEM...  
unclean shutdown of service instance with pid 27511.  
service startup...  
accepting requests at [redacted]  
assigning to volume 3 ...  
run as transaction master  
resetting of user SYSTEM - new password:  
NewPassword1  
new pw accepted.  
(Re)Activating user SYSTEM...  
done
```

#### Reset SYSTEM User Password

7. 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 is reset and the index server stops.

8. In the terminals in which they are running, end the name server and compile server processes by pressing **CTRL+C**.

9. 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 SYSTEM user's password is reset. You do not have to change this new password the next time you log on with this user regardless of your password policy configuration.

If you previously deactivated the SYSTEM user, it is now also reactivated. This means you will need to deactivate it again.

## Related Information

[Configure the Password Policy and Blacklist in SAP HANA Studio \[page 652\]](#)

### 6.2.1.4 Resetting the SYSTEM User Password in Multitenant Database Containers

The system database and all tenant databases in a multiple-container system each have their own SYSTEM user. You can reset the password of these SYSTEM users as the operating system administrator by starting the name server (system database) or index server (tenant database) in emergency mode.

#### 6.2.1.4.1 Reset the SYSTEM User Password of the System Database

If the password of the SYSTEM user in 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.

### **i** 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

1. 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/ C/ /exe/hdbnameserver -resetUserSystem
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.

### 6.2.1.4.2 Reset the SYSTEM User Password of a Tenant Database

If the password of the SYSTEM user in a tenant database is lost, you can reset it as the operating system administrator by starting the index server in emergency mode.

#### Prerequisites

- You cannot log on to the database as the SYSTEM because the password has been irretrievably lost.

#### **i** 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 system privilege DATABASE ADMIN.
- You are logged on as the system administrator user `<sid>adm`.

#### Procedure

- Copy the keys used for internal communication to a temporary location:

```
cp $SECUDIR/sapsrv_internal_<database_name>.pse /tmp
```

- Stop the tenant database by executing the following statement, for example in the SQL console of the SAP HANA studio:

```
ALTER SYSTEM STOP DATABASE <database_name>
```

- Move the temporary keys used for internal communication to the \$SECUDIR directory:

```
mv /tmp/sapsrv_internal_<database_name>.pse $SECUDIR
```

- Log on to the server on which the index server of the tenant database is running as the operating system user (that is, `<osuser>adm` user).
- Open a command-line interface.

6. Open a shell to the primary group that is used for running the tenant database.

```
newgrp <osgroup>
```

7. Export the sap system name and the database name, and then start the index server by executing the following commands:

- `/usr/sap/<SID>/HDB<instance> export SAPSYSTEMNAME=<SID>`
- `/usr/sap/<SID>/HDB<instance> export DBNAME=<database_name>`
- `/usr/sap/<SID>/HDB<instance>/hdbenv.sh`
- `DBNAME=<database_name> /usr/sap/<SID>/HDB<instance>/exe/hdbindexserver -port <internal_port> -resetUserSystem`

The following prompt appears: resetting of user SYSTEM - new password

8. Enter a new password for the SYSTEM user.  
The password for the SYSTEM user of the tenant database is reset and the index server stops.
9. Restart the tenant database:

```
ALTER SYSTEM START DATABASE <database_name>
```

## Results

The password of the SYSTEM user of the tenant database.

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.

## 6.2.2 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.

In addition to the SAP HANA database user SYSTEM, the installation process also creates an external operating system user (<sid>adm, for example, spladm or xyzadm).

This operating system user, also referred to as the operating system administrator, simply 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. Within the SAP HANA studio, the operating system user's credentials are required, for example, to start or stop database processes or to execute a recovery.

### ➔ Tip

As a database administrator, you can securely store the credentials of the operating system user for a system in the SAP HANA studio. To do so, open the system's properties and choose the [SAP System Logon](#) page.

The operating system user is not an SAP HANA database user.

## 6.2.3 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 connection, user passwords are always transmitted in encrypted hashed form during the user authentication process, never in plain text. For HTTP connections, HTTPS must be configured. In SSO environments, we recommend using encrypted communication channels for **all** client connections.

#### [User Authentication Mechanisms \[page 648\]](#)

Authentication mechanisms supported in SAP HANA. Mechanisms that are not required can be disabled.

#### [Configuring SAP HANA for User Authentication and Single-Sign On \[page 651\]](#)

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.

#### [Troubleshooting Problems with User Authentication and SSO \[page 666\]](#)

Authentication problems manifest themselves as failed user logon. In many cases, the reason for the failure will not be clear to the user. You need to analyze the database trace to determine the cause of the problem.

### 6.2.3.1 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
Basic authentication (user name and password)	Users accessing the SAP HANA database authenticate themselves by entering their database user name and password.  For more information, see <i>Password Policy</i> and <i>Password Blacklist</i> .	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"> <li>• Directly from ODBC and JDBC database clients within a network (for example, the SAP HANA studio)</li> <li>• Indirectly from front-end applications such as SAP BusinessObjects applications and other SAP HANA databases using Kerberos delegation</li> <li>• Via HTTP/HTTPS access by means of SAP HANA Extended Services (SAP HANA XS), classic model</li> </ul> <p>In this case, Kerberos authentication is enabled with Simple and Protected GSSAPI Negotiation Mechanism (SPNEGO).</p> <p><b>i Note</b></p> <p>A user who connects to the database using an external authentication provider must also have a database user known to the database. SAP HANA maps the external identity to the identity of an internal database user.</p>	Yes
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 service provider to authenticate users accessing via HTTP/HTTPS by means of SAP HANA XS classic.</p> <p><b>i Note</b></p> <p>A user who connects to the database using an external authentication provider must also have a database user known to the database. SAP HANA maps the external identity to the identity of an internal database user.</p>	Yes
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> <p><b>i Note</b></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>	Yes

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 classic, 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><b>i Note</b></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>	Yes for HTTP/HTTPS access to SAP HANA by means of SAP HANA XS 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><b>i Note</b></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>	Yes
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

## 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`.

### **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.

## 6.2.3.1.1 User Authentication in Multitenant Database Containers

All user authentication mechanisms are supported in multitenant database containers.

Separate, database-specific authentication is possible for every certificate-based authentication mechanism (SAML assertions, X.509 certificates, and logon tickets) since it is possible to create different certificate collections for individual purposes directly in every database. However, 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.

### 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 will be able to log on to other databases in the system.

## 6.2.3.2 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 mechanism(s) you are implementing, you must configure SAP HANA accordingly.

### Configuration of Authentication of SAP HANA XS Classic Applications

You use the Web-based administration tools for SAP HANA XS classic to configure security-related aspects of SAP HANA XS classic applications, including authentication (for example, enforced authentication mechanism, trust store configuration and management, and SAML configuration).

For more information about authentication in SAP HANA XS advanced applications, see *SAP HANA Extended Application Services, Advanced Model* in the *SAP HANA Security Guide*.

### Related Information

[SAP HANA XS Administration Tools \[page 1018\]](#)

## 6.2.3.2.1 Configure the Password Policy and Password Blacklist

Passwords for the basic authentication 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 and maintain entries in the password blacklist in line with your organization's security requirements.

### Context

The password policy is defined by parameters in the `password_policy` section of the `indexserver.ini` system properties file. Although you can configure your password policy directly in the `indexserver.ini` file, it is recommended that you use either the *Password Policy and Blacklist* app of the SAP HANA cockpit or the *Security* editor of the SAP HANA studio.

#### **i** Note

The password policy parameters for the system database of a multiple-container system are maintained in the `nameserver.ini` file, not the `indexserver.ini` file.

In addition to configuring the password policy parameters, you can also add words or partial words to the password blacklist. The password blacklist in SAP HANA is implemented with the table `_SYS_PASSWORD_BLACKLIST` in the schema `_SYS_SECURITY`. This table is empty when you create a new instance.

### Related Information

[Configure the Password Policy and Blacklist in SAP HANA Studio \[page 652\]](#)

[Configure the Password Policy and Blacklist in SAP HANA Cockpit \[page 654\]](#)

## 6.2.3.2.1.1 Configure the 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 system 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
<b>Contained in Password</b>	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.
<b>Case-Sensitive</b>	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 users of the system must be created and changed in line with the defined policy.

## Related Information

[Password Policy Configuration Options \[page 655\]](#)

## 6.2.3.2.1.2 Configure the Password Policy and Blacklist in SAP HANA Cockpit

Configure the password policy and password blacklist using the [Authentication](#) app of the SAP HANA cockpit.

### Prerequisites

- You have the role `sap.hana.security.cockpit.roles::MaintainPasswordPolicy`.
- The [Authentication](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the [SAP HANA Security Overview](#) tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

### Procedure

1. Open the SAP HANA cockpit.
2. Open the [Password Policy and Blacklist](#) app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.
3. Click [Edit](#) in the footer bar.
4. 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](#).
5. 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
<b>Contained in Password</b>	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.
<b>Case-Sensitive</b>	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.

6. Click [Save](#) to save the password policy and password blacklist.

### Related Information

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Open SAP HANA Cockpit \[page 23\]](#)

[Password Policy Configuration Options \[page 655\]](#)

### 6.2.3.2.1.3 Password Policy Configuration Options

The *Password Policy and Blacklist* app in the SAP HANA cockpit and the *Security* editor in the SAP HANA studio allow you to view the password policy and to change its default configuration.

The password policy is defined by parameters in the `password_policy` section of the `indexserver.ini` configuration file. The following sections describe these parameters, which correspond to the configuration options available in the *Password Policy and Blacklist* app and the *Security* editor.

#### **i** Note

The password policy parameters for the system database of a multiple-container system are maintained in the `nameserver.ini` file, not the `indexserver.ini` file.

- [Minimum Password Length \[page 655\]](#)
- [Lowercase Letter/Uppercase Letter/Numerical Digit/Special Character Required \[page 656\]](#)
- [Password Change Required on First Logon \[page 657\]](#)
- [Number of Last Used Passwords That Cannot Be Reused \[page 657\]](#)
- [Number of Allowed Failed Logon Attempts \[page 658\]](#)
- [User Lock Time \[page 658\]](#)
- [Minimum Password Lifetime \[page 659\]](#)
- [Maximum Password Lifetime \[page 659\]](#)
- [Lifetime of Initial Password \[page 660\]](#)
- [Maximum Duration of User Inactivity \[page 660\]](#)
- [Notification of Password Expiration \[page 660\]](#)
- [SYSTEM User Lock \[page 661\]](#)
- [Detailed Error Information on Failed Logon \[page 661\]](#)

## 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 Letter/Uppercase Letter/Numerical Digit/Special Character Required

The character types that the password must contain; at least one character of each selected character type is required

Parameter	password_layout
Default Value	Aa1
Additional Information	<p>The following character types are possible:</p> <ul style="list-style-type: none"> <li>• Lowercase letter (a-z)</li> <li>• Uppercase letter (A-Z)</li> <li>• Numerical digits (0-9)</li> <li>• 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.</li> </ul> <p>The default configuration requires passwords to contain at least one uppercase letter, at least one number, and at least one lowercase letter, with special characters being optional.</p> <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></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 style="background-color: #fff9c4; padding: 5px;"> <p><b>⚠ Caution</b></p> <p>The use of passwords enclosed in double quotes (") may cause logon issues depending on the client used. The SAP HANA Studio, for example, supports passwords enclosed in double quotes ("), while the SAP HANA HDBSQL command line tool does not.</p> </div> <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <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 <b>a1A, hQ5, or 9fG</b>. If you want to enforce the use of at least one of each character type including special characters, you specify <b>A1a_ or 2Bg?</b>.</p> </div>
UI Labels	<i>Lowercase Letter/Uppercase Letter/Numerical Digit/Special Character 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 <b>true</b>, 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 <b>false</b>, 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 &lt;user_name&gt; 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"><li>• <code>CREATE USER &lt;user_name&gt; PASSWORD &lt;password&gt; [NO FORCE_FIRST_PASSWORD_CHANGE]</code></li><li>• <code>ALTER USER &lt;user_name&gt; PASSWORD &lt;password&gt; [NO FORCE_FIRST_PASSWORD_CHANGE]</code></li></ul>
UI Label	<i>Password Change Required on First Logon</i>

## 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 <b>0</b> , 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	maximum_invalid_connect_attempts
Default Value	6 (failed logon attempts)
Additional Information	<p>You must enter a value of at least <b>1</b>.</p> <p>A user administrator can reset the number of invalid logon attempts with the following SQL statement: <code>ALTER USER &lt;user_name&gt; 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 INVALID_CONNECT_ATTEMPTS system view containing the following information:</p> <ul style="list-style-type: none"> <li>• The number of invalid logon attempts since the last successful logon</li> <li>• The time of the last successful logon</li> </ul> <p>A user administrator can delete information about invalid logon attempts with the following SQL statement: <code>ALTER USER &lt;user_name&gt; DROP CONNECT ATTEMPTS</code></p> <div style="background-color: #fff9c4; padding: 10px; margin-top: 10px;"> <p><b>→ Recommendation</b></p> <p>Create an audit policy to log activity in the INVALID_CONNECT_ATTEMPTS 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: #fff9c4; padding: 10px; margin-top: 10px;"> <p><b>i Note</b></p> <p>Although this parameter is not valid for the SYSTEM user, the SYSTEM user will still be locked if the parameter <code>password_lock_for_system_user</code> is set to <b>true</b>. If <code>password_lock_for_system_user</code> is set to <b>false</b>, the SYSTEM 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	password_lock_time
Default Value	1440 (minutes)

Additional Information	<p>If you enter the value <b>0</b>, the user is unlocked immediately. This disables the functionality of parameter <code>maximum_invalid_connect_attempts</code>.</p> <p>A user administrator can reset the number of invalid logon attempts and reactivate the user account with the following SQL statement: <code>ALTER USER &lt;user_name&gt; RESET CONNECT ATTEMPTS</code>. It is also possible to reactivate the user in the user editor of the SAP HANA Studio.</p> <p>To lock a user indefinitely, enter the value <b>-1</b>. In the <i>Security</i> editor of the SAP HANA Studio or the <i>Authentication</i> app of the SAP HANA Cockpit, this corresponds to selecting the <i>Lock User Indefinitely</i> checkbox. The user remains locked until reactivated by a user administrator as described above.</p>
UI Label	<i>User Lock Time</i>

## 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 <b>0</b> , the password has no minimum lifetime.
UI Label	<i>Minimum Password Lifetime</i>

## 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 <b>1</b>.</p> <p>A user administrator can exclude users from this password check with the following SQL statement: <code>ALTER USER &lt;user_name&gt; 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 &lt;user_name&gt; ENABLE PASSWORD LIFETIME</code>.</p>
UI Label	<i>Maximum Password Lifetime</i>

## 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	maximum_unused_initial_password_lifetime
Default Value	7 (days)
Additional Information	You must enter a value of at least <b>1</b> .  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.
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	maximum_unused_productive_password_lifetime
Default Value	365 (days)
Additional Information	You must enter a value of at least <b>1</b> .  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.
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	password_expire_warning_time
Default Value	14 (days)

Additional Information	<p>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.</p> <p>If you enter the value <b>0</b>, the user does not receive notification that his or her password is due to expire.</p> <p>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. For more information about how to disable this check, see SAP Note 1991615.</p>
UI Label	<i>Notification of Password Expiration</i>

## SYSTEM User Not Locked

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
UI Label	<i>SYSTEM User Not Locked</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 <b>false</b>, only the information <code>authentication failed</code> is returned.</p> <p>If set to <b>true</b>, the specific reason for failed logon is returned:</p> <ul style="list-style-type: none"> <li>• Invalid user or password</li> <li>• User is locked</li> <li>• Connect try is outside validity period</li> <li>• User is deactivated</li> </ul>
UI Label	<i>Detailed Error Information on Failed Logon</i>

---

## Related Information

[Execute SQL Statements in SAP HANA Studio \[page 65\]](#)

[Create an Audit Policy \[page 727\]](#)

[SAP Note 1991615](#)

### 6.2.3.2.2 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, 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 terminologyThis 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.

### **i** Note

In an SAP HANA system with multitenant database containers, 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

[SAP HANA XS Administration Tools \[page 1018\]](#)

[Create and Authorize a User \[page 704\]](#)

## 6.2.3.2.3 Add an SAML Identity Provider

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**. For more information, see *Managing Client Certificates in the SAP HANA Database*.

#### **i** Note

If you are using a trust store located in the file system for secure client-server communication over JDBC/ODBC, then this trust store can also be used for SAML-based user authentication. For more information, see *Secure Communication Between SAP HANA and ODBC/JDBC Clients* in the *SAP HANA Security Guide*.

- You have imported into the certificate collection (or trust store in the file system) 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.
- 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 studio or SQL, you should 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"><li>Choose  (<i>Import SAML identity provider from certificate file</i>).</li><li>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.</li></ol> <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><b>i Note</b></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>
<b>Manually</b>	<ol style="list-style-type: none"> <li>1. Choose  (<i>Add SAML identity provider</i>)</li> <li>2. Enter the name of the identity provider (in line with the above naming conventions).</li> <li>3. Enter the issuer and subject DNs.</li> </ol>

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 in the SAP HANA Database \[page 758\]](#)

[Maintaining SAML Providers \[page 1050\]](#)

[Create and Authorize a User \[page 704\]](#)

### 6.2.3.2.4 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, you have all the necessary privileges. For more information, see *Managing Client Certificates in the SAP HANA Database*.
- 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)

If necessary, you can change the location of the trust store in the `indexserver.ini` system properties file.

## Procedure

Configure the trust store:

Option	Description
<b>In-database certificate collection</b>	In the database, create a certificate collection with the purpose <b>SAP LOGON</b> . For more information, see <i>Managing Client Certificates in the SAP HANA Database</i> .
<b>File system based</b>	<ol style="list-style-type: none"><li>1. In the <code>indexserver.ini</code> file, change the value of the <code>[authentication] saplogontickettruststore</code> parameter.</li><li>2. Restart the system.</li></ol>

## Related Information

[Managing Client Certificates in the SAP HANA Database \[page 758\]](#)

### 6.2.3.3 Troubleshooting Problems with User Authentication and SSO

Authentication problems manifest themselves as failed user logon. In many cases, the reason for the failure will not be clear to the user. You need to analyze the database trace to determine the cause of the problem.

For security reasons, no information about error conditions are provided to a user directly after a failed logon attempt, since this could be abused by attackers. In case of authentication problems, the affected user must contact the system administrator, who will then analyze the database trace on the server side.

## Tracing for SSO Issues

Logon information is available in the database trace, but by default, it does not log much information. The first step in troubleshooting any SSO logon issue therefore is to increase the trace level of the authentication-related component(s) of the database trace; you can do this in the SAP HANA studio.

For problems with JDBC/ODBC-based logon, increase the trace level of the `authentication` for the **index sever** to `DEBUG`.

For problems with HTTP-based logon via SAP HANA XS classic, increase the trace level of the `authentication`, `xsession`, and `xsauthentication` components for the **xsengine server** to `DEBUG`.

Once you have increased tracing, reproduce the problematic logon. The traces will now contain more descriptive error messages.

### ➔ Remember

After completing troubleshooting, reduce the authentication trace level back to the default.

In some cases, especially for Kerberos and SPNEGO, it is necessary to use other sources of tracing, such as:

- JDBC, ODBC or SQLDB trace
- Windows event log
- Debugger of browser
- SAP Web dispatcher trace
- Network packet sniffer, for example, Wireshark

For more information about tracing in SAP HANA see the section on traces in the *SAP HANA Administration Guide* and SAP Note 2119087.

### ➔ Tip

Guided Answers is a support tool for troubleshooting problems using decision trees. A guided answer is available for tracing SSO issues.

[Kerberos-Related Authentication Issues \[page 668\]](#)

[SAML Authentication \[page 669\]](#)

## Related Information

[Kerberos-Related Authentication Issues \[page 668\]](#)

[SAML Authentication \[page 669\]](#)

[Traces for SSO Troubleshooting \(Guided Answer\)](#) 

[SAP Note 2119087 - How-To: Configuring SAP HANA Traces](#) 

## 6.2.3.3.1 Kerberos-Related Authentication Issues

### GSS API Errors

Kerberos authentication is implemented in the SAP HANA database using the Generic Security Services Application Program Interface (GSS API). Since GSS API is an internet standard (RFC 4121), all Kerberos-related errors are traced under the `authentication` trace component in the following generic way:

```
<SAP HANA DB error text> (<GSS major code>.<GSS minor code> - <GSS major text> <GSS minor text>)
```

GSS API error texts are sometimes difficult to relate to the concrete problem. The following table contains some hints for selected trace messages.

GSS API Error Code	Error Text	Hint	Solution
851968.252963 9142	Minor error text: Key version number for principal in key table is incorrect	The service key table (keytab) in use on the SAP HANA database host does not match the one created on authentication server.	Re-export the keytab file from the authentication server and re-import it into the host's Kerberos installation.
851968.397560 33	SAP HANA database error text: Cannot get keytab entry for host: <FQDN>  Minor error text: No principal in keytab matches desired name	Keytab actually used might be different than expected (default: /etc/krb5.keytab).	Check environment variable KRB5_KTNAME.
851968.252963 9136	HANA DB error text: Cannot get keytab entry for host: <FQDN>  Minor error text: Configuration file does not specify default realm	Kerberos configuration file actually used might be different than expected (default: /etc/krb5.conf).	Check environment variable KRB5_CONFIG.

### Configuration

There are many potential problems setting up a Kerberos infrastructure that are not related to the SAP HANA system in particular, but relevant for any Kerberos-based authentication. For further information, refer to the documentation provided with MIT Kerberos or Microsoft Server/Active Directory.

### Useful SAP Notes

SAP Note	Description
<a href="#">1837331</a> 	HOWTO HANA DB SSO Kerberos/ Active Directory
<a href="#">2354473</a> 	SSO troubleshooting for HANA and Analysis Office (SPNEGO)

SAP Note	Description
<a href="#">1813724</a>	HANA SSO/Kerberos: create keytab and validate conf
<a href="#">2354556</a>	Common errors when executing hdbkrbconf.py

## 6.2.3.3.2 SAML Authentication

### User cannot connect with SAML assertion

If a user cannot connect to SAP HANA with a SAML assertion, the issuer and subject distinguished names (DNs) in the SAML assertion do not match those configured in the identity provider. Investigate which issuer and subject DN were used in the SAML assertion. You will find them in the trace file `indexserver_alert_<hostname>.trc`. Compare these with those configured in the service provider.

### Useful SAP Notes

SAP Note	Description
<a href="#">1766704</a>	How to use Fiddler to collect HTTP and HTTPS traces
<a href="#">2284620</a>	HOW-TO HANA DB SSO SAML and BI Platform 4.1 / AO 2.2

#### ➔ Tip

Guided Answers is a support tool for troubleshooting problems using decision trees. A guided answer is available for SAML authentication with SAP HANA.

## Related Information

[SAML Authentication for Single Sign-On \(Guided Answer\)](#)

## 6.2.4 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, shutdown, and database recovery), the credentials of the SAP operating system user (<sid>adm) are also required.

## Related Information

[Roles \[page 689\]](#)

### 6.2.4.1 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 privilege DATABASE ADMIN. This system privilege can only be granted to users of the system database. It authorizes 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>

Privilege Type	Applicable To	Target User	Description
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>
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>

Privilege Type	Applicable To	Target User	Description
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 style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></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 style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></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**

In the SAP HANA studio, an additional privilege type can be granted. 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 *Debug an External Session* in the *SAP HANA Developer Guide* .

## 6.2.4.2 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)

In a system with multitenant database containers, system privileges granted to users in a particular database container authorize operations in that database only. The only exception is the system privilege DATABASE ADMIN. This system privilege can only be granted to users of the system database. It authorizes 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.

For more information about the individual system privileges available, see *System Privileges (Reference)*.

### Related Information

[System Privileges \(Reference\) \[page 673\]](#)

## 6.2.4.2.1 System Privileges (Reference)

System privileges control general system activities.

### General System Privileges

System privileges are used to 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 commands: CREATE ADAPTER, DROP ADAPTER and ALTER ADAPTER. It also allows access to ADAPTERS and ADAPTER_LOCATIONS system views.
AGENT ADMIN	Controls the execution of the following agent-related commands: CREATE AGENT, DROP AGENT, and ALTER AGENT. It also allows access to AGENTS and ADAPTER_LOCATIONS system views.
AUDIT ADMIN	Controls the execution of the following auditing-related commands: CREATE AUDIT POLICY, DROP AUDIT POLICY and ALTER AUDIT POLICY and the changes of the auditing configuration. It also allows access to AUDIT_LOG system view.
AUDIT OPERATOR	Authorizes the execution of the following command: ALTER SYSTEM CLEAR AUDIT LOG. It also allows access to AUDIT_LOG system view.
BACKUP ADMIN	Authorizes BACKUP and RECOVERY commands for defining and initiating backup and recovery procedures. It also authorizes changing of system configuration options with respect to backup and recovery.
BACKUP OPERATOR	Authorizes the BACKUP command to initiate a backup.
CATALOG READ	Authorizes users to have unfiltered read-only access to all system views.  Normally, the content of these views is filtered based on the privileges of the accessing user.
CERTIFICATE ADMIN	Authorizes the changing of certificates and certificate collections that are stored in the database.
CREATE R SCRIPT	Authorizes the creation of a procedure using the language R.
CREATE REMOTE SOURCE	Authorizes the creation of remote data sources using the CREATE REMOTE SOURCE command.
CREATE SCENARIO	Controls the creation of calculation scenarios and cubes (calculation database).
CREATE SCHEMA	Authorizes the creation of database schemas using the CREATE SCHEMA command.  By default each user owns one schema, with this privilege the user is allowed to create additional schemas.

System Privilege	Description
CREATE STRUCTURED PRIVILEGE	<p>Authorizes the creation of Structured Privileges (Analytical Privileges).</p> <p>Only the owner of an Analytical Privilege can further grant or revoke that privilege to other users or roles.</p>
CREDENTIAL ADMIN	<p>Authorizes the credential commands: CREATE/ALTER/DROP CREDENTIAL.</p>
DATA ADMIN	<p>Authorizes reading all data in the system views. It also enables execution of any Data Definition Language (DDL) commands in the SAP HANA database.</p> <p>A user with this privilege cannot select or change data stored tables for which they do not have access privileges, but they can drop tables or modify table definitions.</p>
DATABASE ADMIN	<p>Authorizes all commands related to tenant databases, such as CREATE, DROP, ALTER, RENAME, BACKUP, and RECOVERY.</p>
EXPORT	<p>Authorizes export activity in the database via the EXPORT TABLE command.</p> <p>Beside this privilege, the user requires the SELECT privilege on the source tables to be exported.</p>
EXTENDED STORAGE ADMIN	<p>Required to manage SAP HANA dynamic tiering and create extended storage.</p>
IMPORT	<p>Authorizes the import activity in the database using the IMPORT commands.</p> <p>Beside this privilege, the user requires the INSERT privilege on the target tables to be imported.</p>
INIFILE ADMIN	<p>Authorizes changing of system settings.</p>
LICENSE ADMIN	<p>Authorizes the SET SYSTEM LICENSE command to install a new license.</p>
LOG ADMIN	<p>Authorizes the ALTER SYSTEM LOGGING [ON OFF] commands to enable or disable the log flush mechanism.</p>
MONITOR ADMIN	<p>Authorizes the ALTER SYSTEM commands for events.</p>
OPTIMIZER ADMIN	<p>Authorizes the ALTER SYSTEM commands concerning SQL PLAN CACHE and ALTER SYSTEM UPDATE STATISTICS commands, which influence the behavior of the query optimizer.</p>

System Privilege	Description
RESOURCE ADMIN	This privilege authorizes commands concerning system resources, for example ALTER SYSTEM RECLAIM DATAVOLUME and ALTER SYSTEM RESET MONITORING VIEW. It also authorizes many of the commands available in the Management Console.
ROLE ADMIN	<p>This privilege authorizes the creation and deletion of roles using the CREATE ROLE and DROP ROLE commands. It also authorizes the granting and revocation of roles using the GRANT and REVOKE commands.</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 save point process using the ALTER SYSTEM SAVEPOINT command.
SCENARIO ADMIN	Authorizes all calculation scenario-related activities (including creation).
SERVICE ADMIN	<p>Authorizes the ALTER SYSTEM [START CANCEL RECONFIGURE] commands.</p> <p>This privilege is for administering system services of the database</p>
SESSION ADMIN	Authorizes the ALTER SYSTEM commands concerning sessions to stop or disconnect a user session or to change session variables.
SSL ADMIN	Controls the execution of the following commands: SET pse_store_name PURPOSE SSL. It also allows access to the PSES system view.
STRUCTUREDPRIVILEGE ADMIN	Authorizes the creation, reactivation, and dropping of structured privileges.
TENANT ADMIN	Authorizes the tenant operations performed by the ALTER SYSTEM [RESUME SUSPEND] TENANT commands.
TABLE ADMIN	Authorizes the LOAD/UNLOAD/MERGE of tables and its table placement.
TRACE ADMIN	Authorizes the ALTER SYSTEM [CLEAR REMOVE] TRACES commands for operations on database trace files and authorizes changing trace system settings.

System Privilege	Description
TRUST ADMIN	Authorizes commands to update the trust store.
USER ADMIN	Authorizes the creation and modification of users using the CREATE USER, ALTER USER, and DROP USER commands.
VERSION ADMIN	Authorizes the ALTER SYSTEM RECLAIM VERSION SPACE command of the multi-version concurrency control (MVCC) mechanism.
WORKLOAD ADMIN	Authorizes execution of the workload class and mapping commands: CREATE WORKLOAD CLASS, ALTER WORKLOAD CLASS, DROP WORKLOAD CLASS, CREATE WORKLOAD MAPPING, ALTER WORKLOAD MAPPING, and DROP WORKLOAD MAPPING
WORKLOAD ANALYZE ADMIN	Used by Analyze Workload, Capture Workload, and Replay Workload apps when performing workload analysis.
WORKLOAD CAPTURE ADMIN	Authorizes access to monitoring view M_WORKLOAD_CAPTURES to see the current status of capturing and captured workloads, as well of execution of actions with built-in procedure WORKLOAD_CAPTURE
WORKLOAD REPLAY ADMIN	Authorizes access to monitoring views M_WORKLOAD_REPLAY_PREPROCESSES and M_WORKLOAD_REPLAYS to see current status of preprocessing, preprocessed, replaying, and replayed workloads, as well as execution of actions with the built-in procedure WORKLOAD_REPLAY
<identifier>.<identifier>	Components of the SAP HANA database can create new system privileges. These privileges use the component-name as first identifier of the system privilege and the component-privilege-name as the second identifier.

### **i** Note

Additional system privileges (shown as <identifier>.<identifier> above) may exist and be required in conjunction with SAP HANA options and capabilities such as SAP HANA smart data integration. For more information, see *SAP HANA Options and Capabilities* on SAP Help Portal.

## 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
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	

### 6.2.4.3 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.

## 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.

### 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.

For more information about resolving authorization errors with the authorization dependency viewer, see *Resolve Errors Using the Authorization Dependency Viewer* in the *SAP HANA Administration Guide* .

For more information about the object privileges available in SAP HANA and for which objects they are relevant, see *Object Privileges (Reference)*.

## Related Information

[Resolve Errors Using the Authorization Dependency Viewer \[page 694\]](#)

[Object Privileges \(Reference\) \[page 679\]](#)

### 6.2.4.3.1 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 a HANA database.

Object Privilege	Command Types	Applies to	Privilege Description
ALL PRIVILEGES	DDL & DML	<ul style="list-style-type: none"> <li>• Tables</li> <li>• Views</li> </ul>	<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"> <li>• Schemas</li> <li>• Tables</li> <li>• Views</li> <li>• Functions/procedures</li> </ul>	Authorizes the ALTER command for the object.
CREATE ANY	DDL	<ul style="list-style-type: none"> <li>• Schemas</li> </ul>	Authorizes all CREATE commands for the object.
CREATE VIRTUAL FUNCTION	DDL	<ul style="list-style-type: none"> <li>• Remote sources</li> </ul>	Authorizes creation of virtual functions (REFERENCES privilege is also required)
CREATE VIRTUAL FUNCTION PACKAGE	DDL	<ul style="list-style-type: none"> <li>• Schemas</li> </ul>	Authorizes creation of virtual function packages.
CREATE VIRTUAL TABLE	DDL	<ul style="list-style-type: none"> <li>• Remote sources</li> </ul>	Authorizes the creation of proxy tables pointing to remote tables from the source entry
CREATE TEMPORARY TABLE	DDL	<ul style="list-style-type: none"> <li>• Schemas</li> </ul>	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"> <li>• Schemas</li> <li>• Calculation Views</li> <li>• Functions/procedures</li> </ul>	Authorizes debug-functionality for the procedure or calculation view or for the procedures and calculation views of a schema.

Object Privilege	Command Types	Applies to	Privilege Description
DELETE	DML	<ul style="list-style-type: none"> <li>• Schemas</li> <li>• Tables</li> <li>• Views</li> <li>• Functions/procedures</li> </ul>	<p>Authorizes the DELETE and TRUNCATE commands 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"> <li>• Schemas</li> <li>• Tables</li> <li>• Views</li> <li>• Sequences</li> <li>• Functions/procedures</li> <li>• Remote sources</li> </ul>	<p>Authorizes the DROP commands for the object.</p>
EXECUTE	DML	<ul style="list-style-type: none"> <li>• Schemas</li> <li>• Functions/procedures</li> </ul>	<p>Authorizes the execution of an SQLScript function or a database procedure using the CALLS or CALL command respectively. It also allows a user to execute a virtual function.</p>
INDEX	DDL	<ul style="list-style-type: none"> <li>• Schemas</li> <li>• Tables</li> </ul>	<p>Authorizes the creation, modification or dropping of indexes for the object</p>
INSERT	DML	<ul style="list-style-type: none"> <li>• Schemas</li> <li>• Tables</li> <li>• Views</li> </ul>	<p>Authorizes the INSERT command for the object.</p> <p>The INSERT and UPDATE privilege are both required on the object to allow the REPLACE and UPSERT commands to be used.</p> <p>While INSERT 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>

Object Privilege	Command Types	Applies to	Privilege Description
REFERENCES	DDL	<ul style="list-style-type: none"> <li>• Schemas</li> <li>• Tables</li> </ul>	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) for a certain purpose. It also allows a user to reference a virtual function package.
SELECT	DML	<ul style="list-style-type: none"> <li>• Schemas</li> <li>• Tables</li> <li>• Views</li> <li>• Sequences</li> </ul>	Authorizes the SELECT command for this object or the usage of a sequence.
SELECT CDS METADATA	DML	<ul style="list-style-type: none"> <li>• Schemas</li> <li>• Tables</li> </ul>	Authorizes access to CDS metadata from the catalog
SELECT METADATA	DML	<ul style="list-style-type: none"> <li>• Schemas</li> <li>• Tables</li> </ul>	Authorizes access to the complete metadata of all objects in a schema (including procedure and view definitions), thus showing the existence of objects that may be located in other schemas.
TRIGGER	DDL	<ul style="list-style-type: none"> <li>• Schemas</li> <li>• Tables</li> </ul>	Authorizes the CREATE TRIGGER/DROP TRIGGER command for the specified table or the tables in the specified schema.
UPDATE	DML	<ul style="list-style-type: none"> <li>• Schemas</li> <li>• Tables</li> <li>• Views</li> </ul>	<p>Authorizes the UPDATE/LOAD/UNLOAD/LOCK TABLE command for that object.</p> <p>While UPDATE 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>

Object Privilege	Command Types	Applies to	Privilege Description
<identifier>.<identifier>	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.

### Note

Additional object privileges (shown as <identifier>.<identifier> above) may exist and be required in conjunction with SAP HANA options and capabilities such as SAP HANA smart data integration. For more information, see *SAP HANA Options and Capabilities* on SAP Help Portal.

## 6.2.4.4 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).

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"><li>• Attribute views</li><li>• Analytic views</li><li>• Calculation views</li></ul>	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
<b>i</b> Note This corresponds to development in an SAP HANA XS classic environment using the SAP HANA repository.		
Design-time modeling using the SAP Web IDE for SAP HANA	Yes	No
<b>i</b> Note This corresponds to development in an SAP HANA XS advanced environment using HDI.		

Feature	SQL-Based Analytic Privileges	XML-Based Analytic Privileges
Transportable	Yes	Yes
HDI support	Yes	No
Complex filtering	Yes	No

## 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 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.

## 6.2.4.5 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 in the repository, 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:

### **i** 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

Package Privilege	Description
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
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

## 6.2.4.6 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.

### **i** 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"
  ]
}
```

## 6.2.4.7 Roles

A 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 `.hdbrole`).

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).

#### **i** Note

Due to the container-based model of HDI, where each container corresponds to a database schema, roles are schema specific.

In SAP HANA XS advanced applications, database roles control access to database objects only (for example, tables, views, and procedures). Application roles and role collections are used to control and define access to applications. For more information about the authorization concept of XS advanced, see the section *Authorization in SAP HANA XS Advanced* in 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

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.

### → Tip

For more information about security, see the *SAP HANA Security Guide* on the SAP Help Portal.

## 6.2.4.7.1 Catalog Roles and Repository Roles Compared

It is possible to create roles as pure runtime objects that follow classic SQL principles or as design-time objects in the repository of the SAP HANA database. In general, repository roles are recommended as they offer more flexibility. For example, they can be transported between systems.

The following table summarizes the differences between catalog roles and repository roles:

Feature	Catalog Roles	Repository Roles
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"><li>• SAP HANA Application Lifecycle Manager</li><li>• The change and transport system (CTS+) of the SAP NetWeaver ABAP application server</li><li>• SAP HANA Transport Container (HTC)</li></ul>
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.

Feature	Catalog Roles	Repository Roles
Relationship to creating database user	Roles are owned by the database user who creates them. To grant privileges to a role, a user requires all the privileges being granted to the role. If any of these privileges are revoked from the granting user, they are automatically revoked from the role. 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 roles, 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.
Grant and revoke process	Roles created in runtime are granted directly by the database user using the SQL GRANT and REVOKE statements. Roles can only be revoked by the grantor. If the granting user is dropped (not necessarily the role creator), all roles that he or she granted are revoked.	Roles are granted and revoked using built-in procedures. Any administrator with the EXECUTE privilege on these can grant and revoke roles. Role creation is decoupled from the grant and revoke process.

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 the technical user `_SYS_REPO` 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 only be revoked by the same user. 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.

## 6.2.4.8 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.

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 = '&lt;user_name&gt;';</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='&lt;user_name&gt;;</pre>	All application privileges granted to the specified user both directly and indirectly through roles are returned separately.
EFFECTIVE_PRIVILEGE_GRANTEES	<pre>Object privilege: SELECT * FROM EFFECTIVE_PRIVILEGE_GRANTEES WHERE OBJECT_TYPE = '&lt;object_type&gt;' AND SCHEMA_NAME = '&lt;schema_name&gt;' AND OBJECT_NAME = '&lt;object_name&gt;' AND PRIVILEGE = '&lt;privilege&gt;';  System privilege: SELECT * FROM EFFECTIVE_PRIVILEGE_GRANTEES WHERE OBJECT_TYPE = 'SYSTEMPRIVILEGE' AND PRIVILEGE = '&lt;privilege&gt;';  XML-based analytic privilege: SELECT * FROM EFFECTIVE_PRIVILEGE_GRANTEES WHERE OBJECT_TYPE = 'ANALYTICPRIVILEGE' AND SCHEMA_NAME = '&lt;schema_name&gt;' AND OBJECT_NAME = '&lt;privilege_object_name&gt;' AND PRIVILEGE = 'EXECUTE';  SQL-based analytic privilege: SELECT * FROM EFFECTIVE_PRIVILEGE_GRANTEES WHERE OBJECT_TYPE = 'SQLANALYTICPRIVILEGE' AND SCHEMA_NAME = '&lt;schema_name&gt;' AND OBJECT_NAME = '&lt;privilege_object_name&gt;' AND PRIVILEGE = 'EXECUTE';</pre> <div style="background-color: #fff9c4; padding: 10px; margin-top: 10px;"> <p><b>i Note</b></p> <p>For analytic privileges, schema name is optional.</p> </div>	All users or roles who have the specified privilege are returned.

System View	Query	Result
EFFECTIVE_PRIVILEGES	<pre>SELECT * FROM "PUBLIC"."EFFECTIVE_PRIVILEGES" where USER_NAME = '&lt;user_name&gt;;'</pre>	All privileges granted to the specified user both directly and indirectly through roles are returned separately.
EFFECTIVE_ROLE_GRANTEES	<pre>SELECT * FROM EFFECTIVE_ROLE_GRANTEES WHERE ROLE_NAME = '&lt;role_name&gt;;'</pre>	All users or roles who have the specified role are returned.
EFFECTIVE_ROLES	<pre>SELECT * FROM "PUBLIC"."EFFECTIVE_ROLES" where USER_NAME = '&lt;user_name&gt;' AND ROLE_SCHEMA_NAME = '&lt;schema_name of role&gt;;'</pre> <div style="background-color: #fff9c4; padding: 5px; margin-top: 10px;"> <p><b>i Note</b> Schema name is optional.</p> </div>	All roles granted to the specified user both directly and indirectly through other roles are returned separately.
EFFECTIVE_STRUCTURED_PRIVILEGES	<pre>SELECT * from "PUBLIC"."EFFECTIVE_STRUCTURED_P RIVILEGES" where ROOT_SCHEMA_NAME = '&lt;schema&gt;' AND ROOT_OBJECT_NAME = '&lt;object_name&gt;' AND USER_NAME = '&lt;user_name&gt;'</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 = '&lt;user_name&gt;;'</pre>	Privileges granted directly to the specified user (or role) are returned. Privileges contained within granted roles are not shown.  <div style="background-color: #fff9c4; padding: 5px; margin-top: 10px;"> <p><b>i Note</b> It is possible to query the privileges directly granted to a role by replacing where GRANTEE = '&lt;USER&gt;' with where GRANTEE = '&lt;ROLE&gt;'</p> </div>
GRANTED_ROLES	<pre>SELECT * FROM "PUBLIC"."GRANTED_ROLES" where GRANTEE = '&lt;user/role_name&gt;;'</pre>	All roles granted directly to the specified user (or role) are returned. Roles contained within granted roles are not shown.  <div style="background-color: #fff9c4; padding: 5px; margin-top: 10px;"> <p><b>i Note</b> It is possible to query the roles directly granted to a role by replacing where GRANTEE = '&lt;USER&gt;' with where GRANTEE = '&lt;ROLE&gt;'</p> </div>

---

## 6.2.4.9 Resolve Errors Using the Authorization Dependency Viewer

You can use the authorization dependency viewer 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 or DATA ADMIN.

### 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.

### 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 \[page 700\]](#)

---

## 6.2.4.9.1 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 DEPVIEW.
- PROC\_TO\_DEPVIEWER selects from DEPTABLE and DEPVIEW.
- DEPVIEW selects from DEPTABLE.

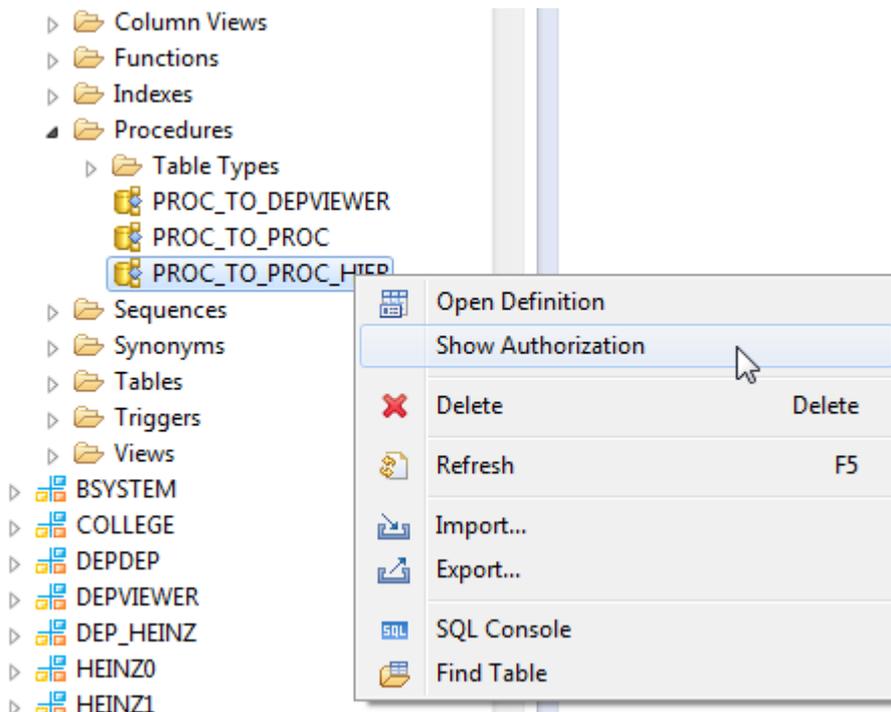
Other users are now granted EXECUTE privilege on PROC\_TO\_PROC\_HIER. However, when they execute the procedure, the following error appears:

```
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)
```

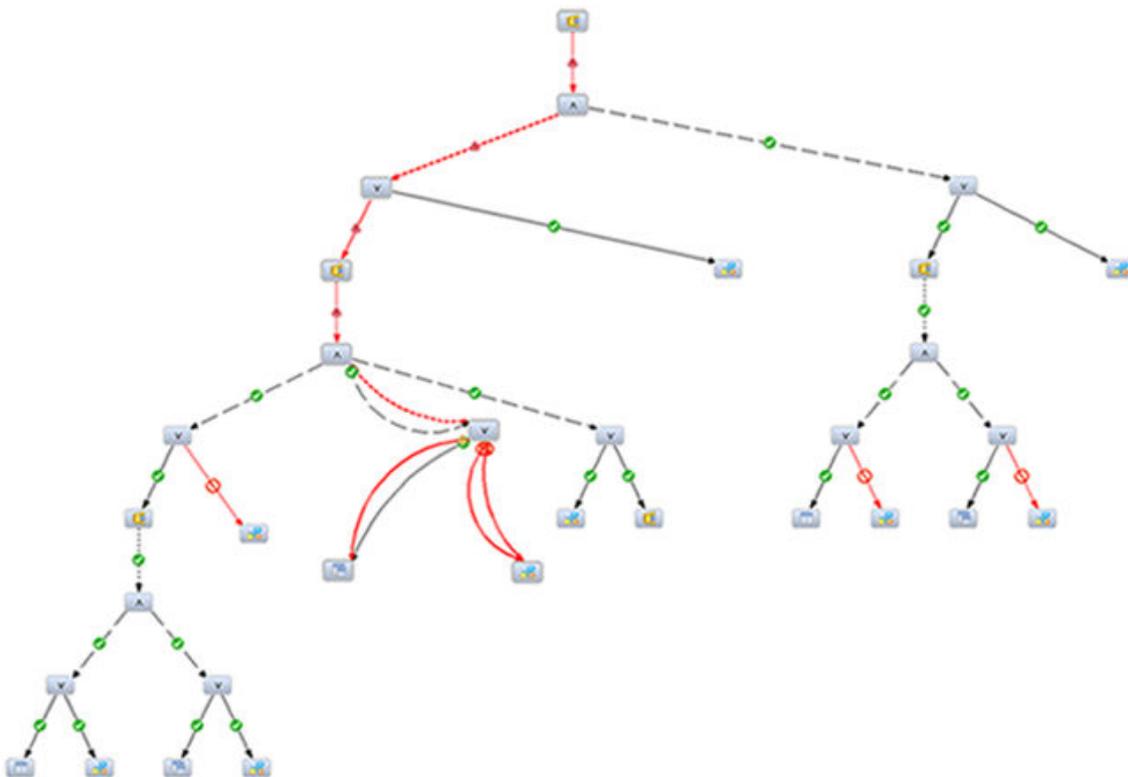
You can use the authorization dependency viewer to isolate the source of the problem as follows:

### Procedure

1. In the *Systems* view, navigate to the procedure PROC\_TO\_PROC\_HIER and from the context menu, choose *Show Authorization*:

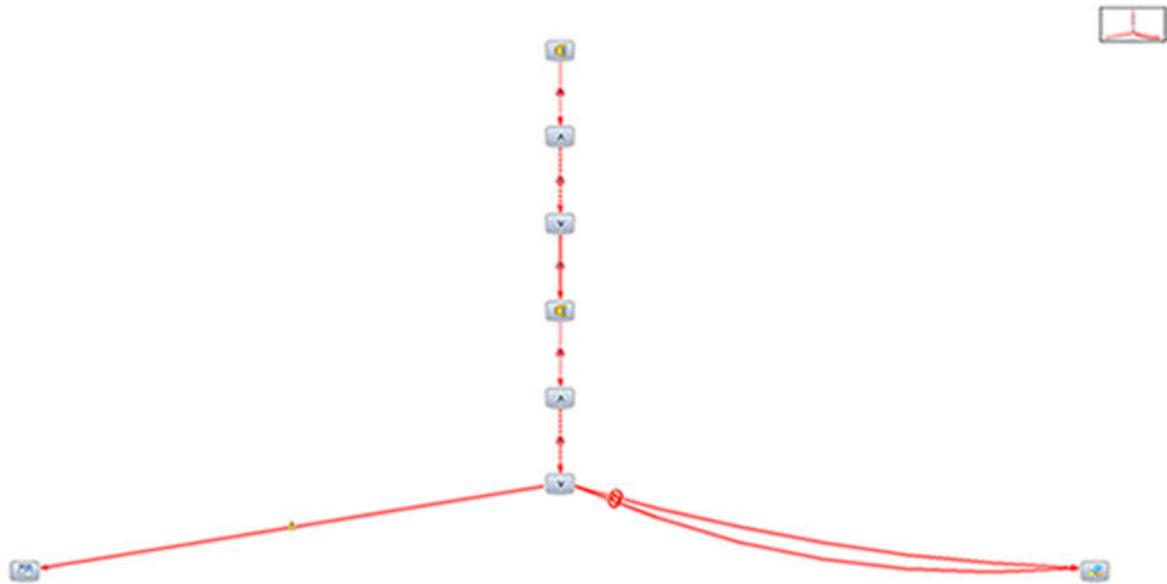


The full authorization dependency structure of the procedure is displayed as a hierarchical tree:



2. 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 DEPVIEWER:

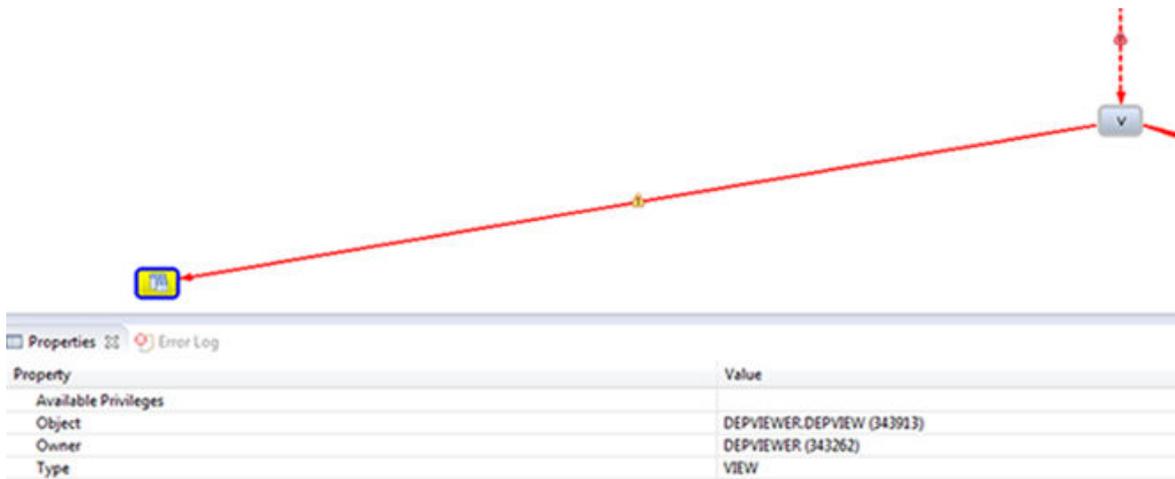


3. To examine the invalid dependency path in more detail, select the connection to the view.

Property	Value
Authorized User	BODOS (138940)
Child Node	DEPVIEW (343913)
Connection Type	OR (1)
Dependency Status	AUTHORIZED NON GRANTABLE
Parent Node	343931_2 (59539)
Required Privilege	DELETE (26)

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.

4. 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.

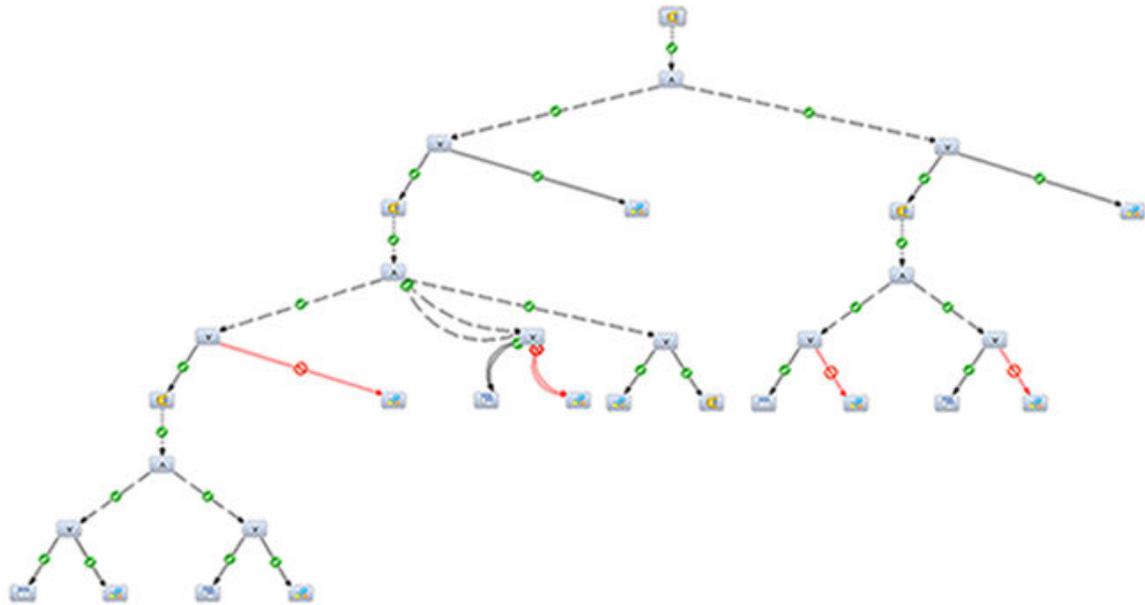
- As user DEPVIEWER, in the user definition of user BODOS, select *Grantable to others* for the EXECUTE privilege on the object DEPVIEW:



### 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.

- In the authorization dependency viewer, choose .  
There are now no invalid authorization dependencies; the procedure is valid ():



## 6.2.4.9.2 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"> <li>  (NOT AUTHORIZED)            The user does not have the required privilege for the child object.         </li> <li>  (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.         </li> <li>  (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.         </li> <li>  (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.         </li> </ul>

### 6.2.4.9.3 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

## 6.2.5 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.

## Related Information

[Catalog Roles and Repository Roles Compared \[page 690\]](#)

## 6.2.5.1 Provisioning Users in SAP HANA Studio

You can use the *User* and *Role* editors of the SAP HANA studio to perform user-provisioning tasks.

### 6.2.5.1.1 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.

#### 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 ("...").
3. 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.

4. 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
5. 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.

For more information, see *SAP HANA Developer Guide* on SAP Help Portal.

## Related Information

[Prerequisites for Granting and Revoking Privileges and Roles \[page 713\]](#)

### 6.2.5.1.2 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.

## 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 user's 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
<b>Authenti-cation</b>	<p>You can set up one or more of the following types of user authentication:</p> <ul style="list-style-type: none"> <li>○ <b>User name/password</b> authentication by specifying a user name and 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.</li> <li>○ <b>Kerberos</b> 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</li> <li>○ <b>SAML</b> 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.</li> <li>○ <b>X.509 certificates</b> by adding the user's public key certificate(s)</li> </ul> <div style="background-color: #fff9c4; padding: 5px; margin: 5px 0;"> <p><b>i</b> Note</p> <p>X.509 certificates are supported only for HTTP access through SAP HANA XS</p> </div> <ul style="list-style-type: none"> <li>○ SAP logon and assertion tickets</li> </ul>
<b>Valid From/Until</b>	<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>
<b>Session client</b>	<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 *Additional User Parameters*.

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.

For more information, see also *SAP HANA SQL and System Views Reference* on SAP Help Portal.

## Related Information

[Prerequisites for Granting and Revoking Privileges and Roles \[page 713\]](#)

[Database Users \[page 638\]](#)

[Configure the Password Policy and Blacklist in SAP HANA Studio \[page 652\]](#)

[Password Policy Configuration Options \[page 655\]](#)

[Configure Kerberos for SAP HANA Database Hosts \[page 662\]](#)

[Add an SAML Identity Provider \[page 664\]](#)

[Maintaining Single Sign-On for SAP HANA XS Applications \[page 1074\]](#)

[Monitoring Memory Usage \[page 270\]](#)

[Additional User Parameters \[page 716\]](#)

### 6.2.5.1.3 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.

#### 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 user's 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
<b>Authenti- cation</b>	<p>You can set up one or more of the following types of user authentication:</p> <ul style="list-style-type: none"> <li>○ <b>User name/password</b> authentication by specifying a user name and 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.</li> <li>○ <b>Kerberos</b> 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</li> <li>○ <b>SAML</b> 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.</li> <li>○ <b>X.509 certificates</b> by adding the user's public key certificate(s)</li> </ul> <div style="background-color: #fff9c4; padding: 5px; margin: 10px 0;"> <p><b>i Note</b></p> <p>X.509 certificates are supported only for HTTP access through SAP HANA XS</p> </div> <ul style="list-style-type: none"> <li>○ SAP logon and assertion tickets</li> </ul>
<b>Valid From/ Until</b>	<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>

Option	Description
<b>Session client</b>	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.

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 [Additional User Parameters](#).

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](#).

For more information, see also [SAP HANA SQL and System Views Reference](#) on SAP Help Portal.

## Related Information

[Prerequisites for Granting and Revoking Privileges and Roles \[page 713\]](#)

[Database Users \[page 638\]](#)

[Configure the Password Policy and Blacklist in SAP HANA Studio \[page 652\]](#)

[Configure Kerberos for SAP HANA Database Hosts \[page 662\]](#)

[Add an SAML Identity Provider \[page 664\]](#)

[Maintaining Single Sign-On for SAP HANA XS Applications \[page 1074\]](#)

[Monitoring Memory Usage \[page 270\]](#)

[Additional User Parameters \[page 716\]](#)

## 6.2.5.1.4 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.

### 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.

#### 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.

### 6.2.5.1.5 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.

#### Prerequisites

- You have the system privilege USER ADMIN.

#### **i** 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* and also *SAP HANA SQL and System Views Reference* on SAP Help Portal.
- 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 \[page 713\]](#)

[Configure the Password Policy and Blacklist in SAP HANA Studio \[page 652\]](#)

[Configure Kerberos for SAP HANA Database Hosts \[page 662\]](#)

[Add an SAML Identity Provider \[page 664\]](#)

[Maintaining Single Sign-On for SAP HANA XS Applications \[page 1074\]](#)

### 6.2.5.1.6 Delete a User

You may need to delete a database user if an employee leaves your organization for example.

## 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.

## 6.2.5.1.7 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 if a security violation is detected.

### 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.

## Related Information

[SAP Note 1986645](#)

### 6.2.5.1.8 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.

#### 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...**  
You are prompted to enter a new password for the user. The user is now reactivated.

### 6.2.5.1.9 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

To grant...	The granting user needs...
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
A role created in runtime	Either: <ul style="list-style-type: none"> <li>The role being granted and be authorized to grant it to other users and roles, or</li> <li>The system privilege ROLE ADMIN</li> </ul>
A role created in the repository	The object privilege EXECUTE on the procedure GRANT_ACTIVATED_ROLE

#### Prerequisites for Revoking Privileges

To revoke ...	The granting 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
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
A role created in runtime	To be the user who granted the role
A role created in the repository	The object privilege EXECUTE on the procedure REVOKE_ACTIVATED_ROLE

## 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, see *Roles as Repository Objects* in the SAP HANA Security Guide.

## Related Information

[Create and Authorize a User \[page 704\]](#)

[Create and Authorize a Restricted User \[page 706\]](#)

[SAP HANA Security Guide](#)

### 6.2.5.1.10 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+002D	-	Hyphen-minus
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
U+005C	\	Reverse solidus
U+005D	]	Right square bracket
U+005E	^	Circumflex accent

Unicode Character	Character	Name
U+0060	`	Grave accent
U+007B	{	Left curly bracket
U+007C		Vertical line
U+007D	}	Right curly bracket
U+007E	~	Tilde

### 6.2.5.1.11 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.
LOCALE	The user's locale 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.
PRIORITY	The priority with which the thread scheduler handles statements executed by the user The priority can be in the range 0-9 with 9 representing the highest priority. 5 is the default priority.
STATEMENT MEMORY LIMIT	The maximum memory (in GB) that can be used by a statement executed by the user 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.  <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.  With this user parameter, you can set a user-specific limit that takes precedence over the global statement memory limit.  For more information about memory usage, see <i>Monitoring Memory Usage</i> .
TIME ZONE	The user's timezone The standard database formats for locale and timezone are supported.

---

## Related Information

[Monitoring Memory Usage \[page 270\]](#)

[Setting a Memory Limit for SQL Statements \[page 277\]](#)

### 6.2.5.2 Provisioning Users in SAP HANA Cockpit

You can use the *Role Assignment* app of the SAP HANA cockpit to grant roles to users.

## Related Information

[SAP HANA Cockpit \[page 22\]](#)

[Open SAP HANA Cockpit \[page 23\]](#)

[Open SAP HANA Cockpit from SAP HANA Studio \[page 25\]](#)

[Assign Roles to a User \[page 717\]](#)

#### 6.2.5.2.1 Assign Roles to a User

Roles are the standard mechanism of granting privileges to users in SAP HANA. It is recommended that you assign roles to users instead of granting privileges individually.

## Prerequisites

- You have the privileges granted by role `sap.hana.security.cockpit.roles::EditAssignedRoles`.
- The *Assign Roles to Users* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA User Management* tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.
- 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)*.

---

## Procedure

1. Open the *Assign Roles* app by clicking the *Assign Roles to Users* tile on the homepage of the SAP HANA cockpit.
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 further roles by clicking *Assign Roles*, selecting the relevant roles and clicking *OK*.
5. Save the user.

The user is granted the selected roles.

## Related Information

[Open SAP HANA Cockpit \[page 23\]](#)

[Open SAP HANA Cockpit from SAP HANA Studio \[page 25\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Roles \[page 689\]](#)

[Create and Authorize a User \[page 704\]](#)

## 6.3 Auditing Activity in SAP HANA Systems

Auditing provides you with visibility on who did what in the SAP HANA database (or tried to do what) and when.

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.

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.

If the audit trail target is a database table, you can avoid the audit table growing indefinitely by deleting audit entries created up until a certain time and date.

You can use the SAP HANA cockpit or the SAP HANA studio to enable auditing, configure audit trail targets, and create audit policies.

---

## Multitenant Database Containers

Auditing can be enabled individually for every database in a multiple-container system. For tenant databases, the underlying system property (`[auditing configuration] global_auditing_state`) is set at the database layer of the `global.ini` file. For the system database, it is set in the `nameserver.ini` file.

Tenant database administrators **cannot** configure audit trail targets independently for their database. The default target for all audit trails in tenant databases is internal database table. The system administrator may change the default audit trail targets for tenant databases by changing the underlying property (`[auditing configuration] *_audit_trail_type`) in the `global.ini` file.

### 6.3.1 Managing Auditing in the SAP HANA Cockpit

Use the *Auditing* app of the SAP HANA cockpit to enable auditing, configure audit trail targets, and create audit policies.

#### [Activate and Configure Auditing \[page 719\]](#)

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 *Auditing* app of the SAP HANA cockpit.

#### [Create an Audit Policy \[page 721\]](#)

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 *Auditing* app of the SAP HANA cockpit.

#### [Auditing Details \[page 724\]](#)

The *Auditing* app of the SAP HANA cockpit allows you to view the details of all audit policies in the SAP HANA database, as well as the configured audit trail targets.

#### 6.3.1.1 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 using the *Auditing* app of the SAP HANA cockpit.

#### Prerequisites

- You have the privileges granted by the role `sap.hana.security.cockpit.roles::MaintainAuditPolicy`.

- The *Auditing* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the [SAP HANA Security Overview](#) tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

## Procedure

1. Open the *Auditing* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.
2. Enable auditing.
  - a. Open the *Configuration* tab and choose *Edit* in the footer bar.
  - b. Set the auditing status to *Enable*.

### **i** Note

In multiple-container systems, you enable auditing for the system database and each tenant database independently.

3. Configure the required audit trail targets.

You can configure multiple audit trail targets: one for the system (*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 audit trail target configured for the system.

*Database table* is the default audit trail target. 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*.

4. Save your configuration.

## Results

Auditing is now activated in your system (or database) and you can create the required audit policies.

**Task overview:** [Managing Auditing in the SAP HANA Cockpit \[page 719\]](#)

## Related Information

[Create an Audit Policy \[page 721\]](#)

## 6.3.1.2 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 using the [Auditing](#) app of the SAP HANA cockpit.

### Prerequisites

- You have the privileges granted by the role `sap.hana.security.cockpit.roles::MaintainAuditPolicy`.
- The [Auditing](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the [SAP HANA Security Overview](#) tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

### Procedure

1. Open the [Auditing](#) app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.
2. On the [Audit Policies](#) tab, 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. 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
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.

#### 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. 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.

7. Optional: Select 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*.

### 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*.

8. Specify the actions to be audited by clicking the add button and selecting the relevant 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.

9. If necessary, specify the target object(s) to be audited by clicking the add 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 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. Alternatively, you can specify that the actions in the policy be audited when performed by all users **except** a particular user or users.

---

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.

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.

**Task overview:** [Managing Auditing in the SAP HANA Cockpit \[page 719\]](#)

## Related Information

[Activate and Configure Auditing \[page 719\]](#)

[Auditing Details \[page 724\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Audit Trail Targets \[page 731\]](#)

## 6.3.1.3 Auditing Details

The *Auditing* app of the SAP HANA cockpit allows you to view the details of all audit policies in the SAP HANA database, as well as the configured audit trail targets.

### Audit Policies

Field	Description
<i>Policy Name</i>	Audit policy name
<i>Policy Status</i>	Audit policy status A policy can be either <i>Enabled</i> or <i>Disabled</i> .
<i>Audited Actions</i>	The action to be audited 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> .
<i>Audited Action Status</i>	When the actions in the policy are to be audited: <ul style="list-style-type: none"><li>• On successful execution</li><li>• On unsuccessful execution</li><li>• On both successful and unsuccessful execution</li></ul>
<i>Audit Level</i>	The severity of the audit entry written to the audit trail when the actions in the policy occur The following audit levels are possible <ul style="list-style-type: none"><li>• Emergency</li><li>• Critical</li><li>• Alert</li><li>• Warning</li><li>• Info</li></ul>
<i>Users</i>	User(s) included in the audit policy or excluded from the audit policy Actions in the policy are audited when performed by either the specified user(s) or any user except the specified user(s).

Field	Description
<i>Target Object</i>	<p>Audited object(s)</p> <p>The following target object types are possible:</p> <ul style="list-style-type: none"> <li>• Schemas (and all objects contained within)</li> <li>• Tables</li> <li>• Views</li> <li>• Procedures</li> <li>• Sequences</li> </ul>
<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 system. The applicable default audit trail target is always indicated in the brackets.</p>

## 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

**Parent topic:** [Managing Auditing in the SAP HANA Cockpit \[page 719\]](#)

## Related Information

[Activate and Configure Auditing \[page 719\]](#)

[Create an Audit Policy \[page 721\]](#)

[Auditing Activity in SAP HANA Systems \[page 718\]](#)

## 6.3.2 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.

### 6.3.2.1 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 \[page 731\]](#)

### 6.3.2.2 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.

### 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.

### 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.

8. 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*.

The following audit trail targets are supported:

9. 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 \[page 731\]](#)

[Best Practices and Recommendations for Creating Audit Policies \[page 732\]](#)

[Editors and Views of the SAP HANA Administration Console \[page 63\]](#)

### 6.3.2.3 Delete Audit Entries from the Audit Trail

If the audit trail target is a database table, you can avoid the audit table growing indefinitely by deleting audit entries created up until a certain time and date.

#### 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.

#### 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 \[page 251\]](#)

[Configure Alerting Thresholds with SAP HANA Cockpit \[page 306\]](#)

### 6.3.3 Audit Trail Targets

In production systems, SAP HANA supports syslog and database table as audit trail targets.

Audit Trail Target	Description
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> <div style="background-color: #fff9c4; padding: 10px;"><p><b>⚠ Caution</b></p><p>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.</p></div>
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 view 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. 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>

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

Administration editor of the SAP HANA studio, it is listed on the *Diagnosis Files* tab, 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).

## 6.3.4 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.

## 6.4 Managing Encryption Keys

SAP HANA generates unique encryption keys on installation for all mechanisms used in SAP HANA to encrypt data. However, if you received SAP HANA pre-installed from a hardware or hosting partner, you might want to change encryption keys to ensure they are not known outside your organization.

### [Server-Side Data Encryption \[page 735\]](#)

SAP HANA features two data encryption services: data encryption in the persistence layer and an internal encryption service available to applications requiring data encryption. SAP HANA uses the secure store in the file system (SSFS) to protect the root keys for these encryption services.

### [Client-Side Data Encryption \(hdbuserstore\) \[page 748\]](#)

The secure user store (`hdbuserstore`) is a tool installed with the SAP HANA client. You 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.

## 6.4.1 Server-Side Data Encryption

SAP HANA features two data encryption services: data encryption in the persistence layer and an internal encryption service available to applications requiring data encryption. SAP HANA uses the secure store in the file system (SSFS) to protect the root keys for these encryption services.

### Data Volume Encryption

What Does It Do?	What Encryption Keys Are Involved?
<p>If enabled, this internal encryption service protects all data saved to disk from unauthorized access at operating system level.</p> <p>For more information, see <i>Data Volume Encryption</i> in the SAP HANA Security Guide.</p>	<p>Pages in the data area are encrypted using page encryption keys. Page encryption keys are encrypted with the data volume encryption root key.</p> <p>In a system that supports multitenant database containers, the system database and all tenant database have their own root key.</p> <p>The root key is generated randomly during installation. The page keys are created when data volume encryption is enabled.</p>

## Internal Data Encryption Service

What Does It Do?	What Encryption Keys Are Involved?
<p>This internal encryption service is used in the following contexts:</p> <ul style="list-style-type: none"> <li> <b>Secure internal credential store</b>            This service stores credentials required by SAP HANA for outbound connections. It is used when data is retrieved from remote data sources using SAP HANA smart data access. It is also used during HTTP destination calls from SAP HANA XS applications.            For more information, see <i>Secure Internal Credential Store</i> in the SAP HANA Security Guide.         </li> <li> <b>Secure stores defined using the SAP HANA XS <code>\$.security.Store</code> API</b>            Application developers can create XS secure stores to store certain application data in name-value form. For more information, see <i>Using the Server-Side JavaScript APIs</i> in the <i>SAP HANA Developer Guide (For SAP HANA Studio)</i> and the <i>Class:Store</i> in the <i>SAP HANA XS JavaScript API Reference</i> .         </li> <li> <b>Private key store</b>            This service stores the private keys of the SAP HANA server required for secure client-server communication, if the relevant personal security environment (PSE) is stored in the database.            For more information, see <i>SSL Configuration on the SAP HANA Server</i> and <i>Certificate Management in SAP HANA</i> in the SAP HANA Security Guide.         </li> </ul>	<p>Every consumer of the service has its own system-internal application encryption key. Application encryption keys are encrypted with the data encryption service root key.</p> <p>In a system that supports multitenant database containers, the system database and all tenant database have their own root key.</p> <p>The root key is generated randomly during installation. The application key for the internal credential store is generated randomly during the first startup. Application keys for XS secure stores are created with the XS secure store. The application key for the private key store is created when the first private key is set for an in-database PSE.</p>

## Instance Secure Store in the File System (SSFS)

What Does It Do?	What Encryption Keys Are Involved?
<p>This secure store stores internal root keys in the file system.</p>	<p>The master key of the instance SSFS encrypts the data volume encryption root key and the data encryption service root key.</p> <div data-bbox="805 539 1396 808" style="background-color: #fff9c4; padding: 10px;"> <p><b>→ Recommendation</b></p> <p>The initial master key that protects the instance SSFS is changed during installation or update. If you received your system pre-installed from a hardware or hosting partner, we recommend that you change it immediately after handover to ensure that it is not known outside of your organization.</p> </div> <div data-bbox="805 819 1396 1037" style="background-color: #fff9c4; padding: 10px;"> <p><b>i Note</b></p> <p>The default path of the key file is <code>\$DIR_GLOBAL/hdb/security/ssfs</code>. If you change the default path, you may need to reconfigure it in the event of a system rename.</p> </div>

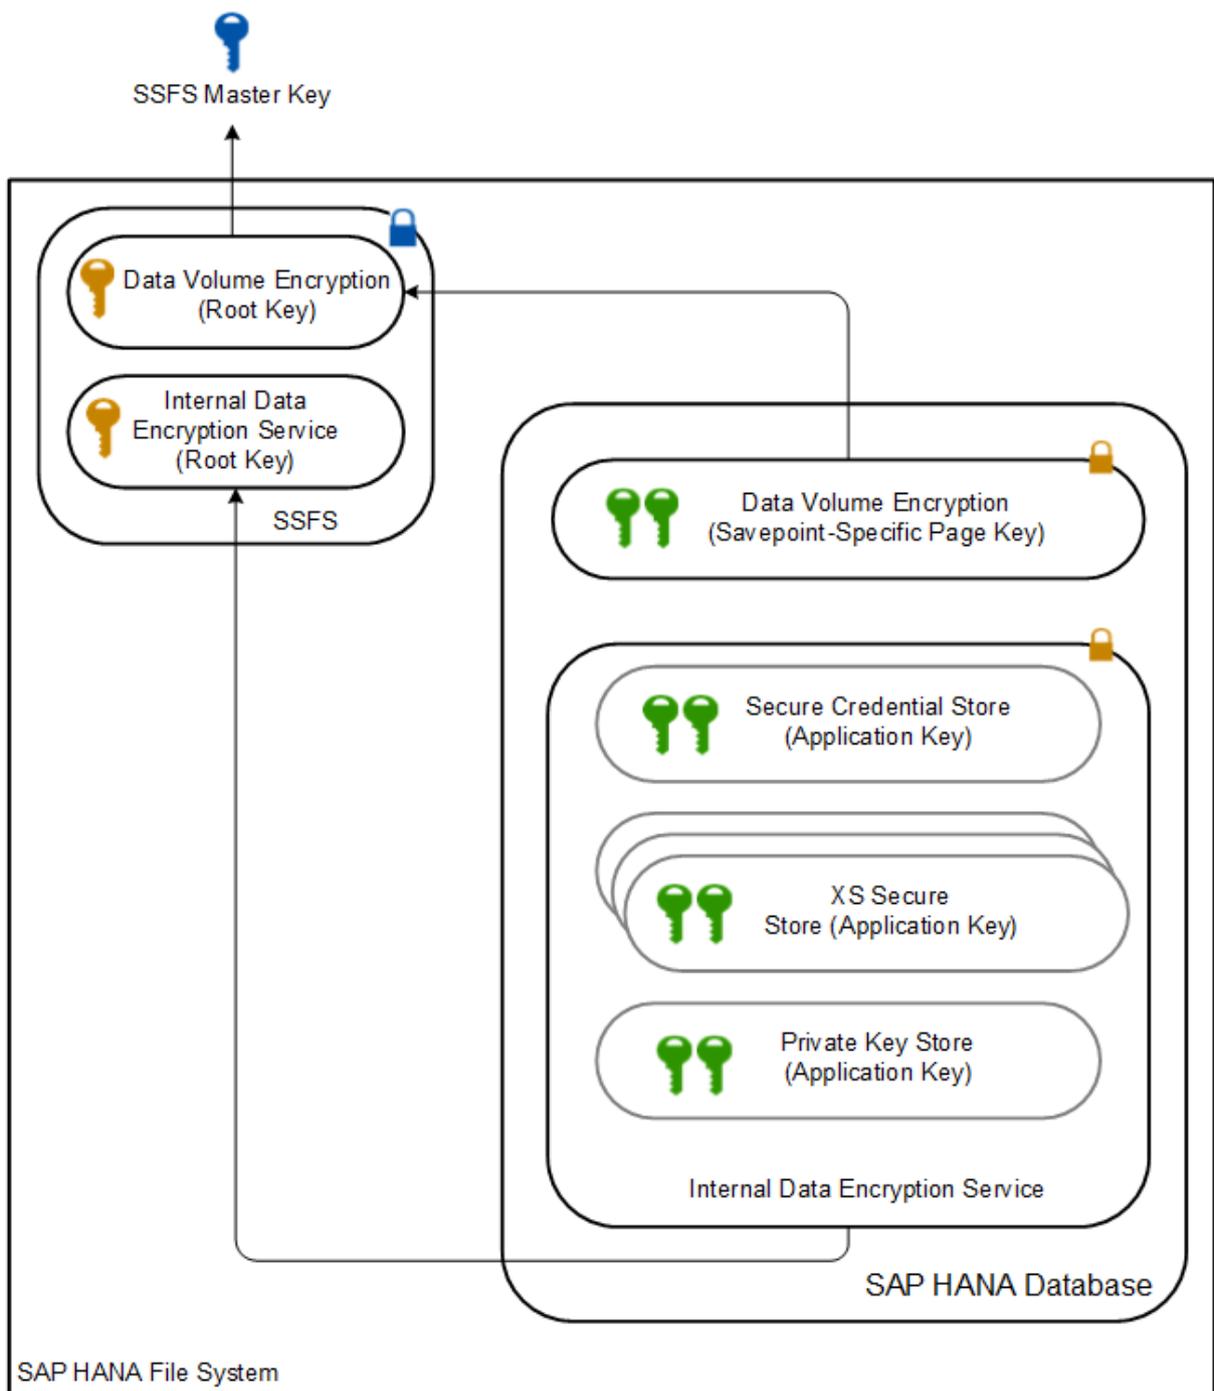
## System PKI (Public Key Infrastructure) SSFS

What Does It Do?	What Encryption Keys Are Involved?
<p>This secure store stores internal root keys in the file system.</p>	<p>The master key of the system PKI SSFS protects the X.509 certificate infrastructure that is used to secure internal SSL/TLS-based communication between hosts in a multiple-host system or between processes of individual databases in a multiple-container system.</p> <div data-bbox="805 1480 1396 1753" style="background-color: #fff9c4; padding: 10px;"> <p><b>→ Recommendation</b></p> <p>The initial master key that protects the system PKI SSFS is changed during installation or upgrade. If you received your system pre-installed from a hardware or hosting partner, we recommend that you change it immediately after handover to ensure that it is not known outside of your organization.</p> </div>

The following figure illustrates the encryption keys protected by the instance SSFS.

### **i Note**

The system PKI SSFS is not depicted. For more information about the system PKI, see *Secure Internal Communication*.



Encryption Keys Protected by the Instance SSFS

## Related Information

[SAP HANA XS JavaScript API Reference](#)

## 6.4.1.1 Encryption Key Management

SAP HANA generates unique root keys on installation. However, if you received SAP HANA pre-installed from a hardware vendor, you might want to change them to ensure they are not known outside your organization. We recommend that you do this immediately after handover from your hardware partner.

The following root keys exist and can be changed:

- Instance SSFS master key
- System PKI SSFS master key
- Data volume encryption root key
- Data encryption service root key

Reinstalling your system will change all master and root keys. You can change keys manually and individually.

The following sections explain how and when you can safely change root keys. More detailed instructions are available in the *SAP HANA Administration Guide*.

### SSFS Master Keys

How to Change	When to Change
<p>Using the command line tool <code>rsecssfx</code></p> <p>The commands are: <code>generatekey</code> and <code>changekey</code></p> <p>➔ <b>Remember</b></p> <p>You'll need operating system access (&lt;sid&gt;adm user) to execute <code>rsecssfx</code> commands.</p> <p>For more information, see <i>Change the SSFS Master Keys</i> in the <i>SAP HANA Administration Guide</i>.</p>	<p>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.</p> <p>You can also change the master keys any time later.</p> <p><b>i Note</b></p> <p>In a system-replication configuration, you change the SSFS master keys on the primary system. To trigger replication of new keys to the secondary system, you must subsequently restart the secondary system. In multitier 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 keys have been replicated, the new keys will be overwritten with the old ones.</p>

## Data Volume Encryption Root Key

How to Change	When to Change
<p>Using the <i>Data Volume Encryption</i> app of the SAP HANA cockpit or the SQL command <code>ALTER SYSTEM PERSISTENCE ENCRYPTION CREATE NEW ROOT KEY</code></p> <p>For more information, see <i>Change the Root Encryption Key for Data Volume Encryption</i> in the <i>SAP HANA Administration Guide</i>.</p>	<p>A unique root key is generated during installation or update. However, if you received your system pre-installed from a hardware or hosting partner, we recommend that you change it immediately after handover to ensure that it is not known outside of your organization.</p> <p>You can also change it any time later.</p> <div data-bbox="807 645 1396 855" style="background-color: #fff9c4; padding: 10px;"><p><b>i Note</b></p><p>In a system-replication configuration, change the root key used for data volume encryption in the primary system only. The new key will be propagated to all secondary systems.</p></div>

## Data Encryption Service Root Key

How to Change	When to Change
<p>Using the command line tool <code>hdbnsutil</code></p> <p>The command is: <code>generateRootKeys type=DPAPI</code></p> <div data-bbox="199 526 785 683"><p>➔ <b>Remember</b></p><p>You'll need operating system access (&lt;sid&gt;adm user) to execute <code>hdbnsutil</code> commands.</p></div> <div data-bbox="199 694 785 1025"><p><b>⚠ Caution</b></p><p>After you change the root key with the command <code>generateRootKeys type=DPAPI</code>, you must <b>immediately</b> do the following two things:</p><ul style="list-style-type: none"><li>• Reset the consistency information in the SSFS using the SAP HANA tool <code>hdbcons</code>.</li><li>• Change all application keys so that they are encrypted with the new root key.</li></ul></div> <p>For more information, see <i>Change the Data Encryption Service Root Key</i> in the <i>SAP HANA Administration Guide</i>.</p>	<p>A unique root key is generated during installation or update. However, if you received your system pre-installed from a hardware or hosting partner, we recommend that you change it immediately after handover to ensure that it is not known outside of your organization.</p> <p>You must change this key at the latest before any data is encrypted using the service. This means before you create any of the following things:</p> <ul style="list-style-type: none"><li>• A remote data source</li><li>• A HTTP destination</li><li>• An XS secure store</li><li>• A certificate collection with a private key</li></ul> <p>You can use the following system views to see whether any data has already been encrypted:</p> <ul style="list-style-type: none"><li>• <code>CREDENTIALS (PUBLIC)</code> If the credential store is empty, then this view will also be empty.</li><li>• <code>P_DPAPI_KEY_ (SYS)</code> If there are no XS secure stores, then this view will have no records with the caller <code>XsEngine</code>. If there are no certificate collections with private keys, there will be no records with the caller <code>PSEStore</code>. Only the user <code>SYSTEM</code> can access this view.</li></ul> <p>Additionally, if the system supports multitenant database containers, you must change the root key before any tenant databases have been created.</p> <div data-bbox="810 1384 1390 1742"><p><b>⚠ Caution</b></p><p>It is important that you plan this root key change carefully as you will have to shut down the database. Not only that, but changing the root key after data has been encrypted will result in key information in the SSFS and the database becoming inconsistent and encrypted data becoming inaccessible. Rectifying the problem could result in data loss. We recommend that you contact SAP Support if errors related to inconsistent SSFS or encryption failure occur.</p></div>

## Related Information

[Change the SSFS Master Keys \[page 742\]](#)

[Change the Root Encryption Key for Data Volume Encryption \[page 744\]](#)

---

[Change the Root Key of the Internal Data Encryption Service \[page 745\]](#)

[SAP Note 2097613](#) 

## 6.4.1.2 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 the credentials of the operating system user (<sid>adm user) 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 root encryption keys listed below. These root keys protect all encryption keys used in the SAP HANA database from unauthorized access.

- The root key used for the internal data encryption service of the database
- The root key used for data volume encryption

In a system that supports multitenant database containers, the system database and all tenant databases have their own root encryption keys for both the data encryption service and data volume encryption.

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 a multiple-container system.

You can change the SSFS master keys using the command line tool `rsecssfsx`, 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.
- In a system that supports multitenant database containers, the SSFS master keys only have to be changed once for whole instance and not per tenant database.
- In a system-replication configuration, you change the SSFS master keys on the primary system. To trigger replication of new keys to the secondary system, you must subsequently restart the secondary system. In multitier 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 keys have been replicated, the new keys will be overwritten with the old ones.

## 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>
```

### **i** Note

The default path of the key file is `$DIR_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 multitier 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 (including snapshots), you don't have to take any actions regarding the master key since only the content of the instance SSFS but not the master key is contained in the backup.

### **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 \[page 98\]](#)

[Change a System Property \[page 217\]](#)

[Server-Side Data Encryption \[page 735\]](#)

[Secure Storage in the File System \(AS ABAP\)](#)

### 6.4.1.3 Change the Root Encryption Key for Data Volume Encryption

SAP HANA generates unique root keys on installation. However, if you received SAP HANA pre-installed from a hardware or hosting partner, you might want to change the root key used for data volume encryption to ensure it is not known outside your organization. You can change the root key using the [Data Volume Encryption](#) app of the SAP HANA cockpit.

#### Prerequisites

- You have the privileges granted by the role `sap.hana.security.cockpit.roles::MaintainDataVolumeEncryption`.
- The [Data Storage Security](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the [SAP HANA Security Overview](#) tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

#### Procedure

##### **i** Note

If the system is configured for system replication, change the root key in the primary system only. The new key will be propagated to all secondary systems.

1. In the SAP HANA cockpit, open the [Data Volume Encryption](#) app.
2. Choose [Change Root Encryption Key](#) in the footer bar.

A root key change request is sent to the server and the value in the [Root Key Change Pending](#) column changes to [Yes](#).

## Results

A new unique root encryption key is generated after the next savepoint operation and all page encryption keys are re-encrypted with the new key. When this process has completed, the value in the *Root Key Change Pending* column changes to *No* and the value in the *Current Key Version* increments by one.

### 6.4.1.4 Change the Root Key of the Internal Data Encryption Service

SAP HANA generates unique root keys on installation. However, if you received SAP HANA pre-installed from a hardware or hosting partner, you might want to change the root key of the internal data encryption service to ensure it is not known outside your organization. We recommend that you do this immediately after system installation or handover from your hardware or hosting partner.

#### Prerequisites

- You have the credentials of the operating system user (<sid>adm user) that was created when the system was installed.
- You have the credentials of the database user SYSTEM.
- You have the system privilege RESOURCE ADMIN.

#### Context

The internal data encryption service can be used by SAP HANA XS-based applications and SAP HANA internal components to securely store data in the database. Consumers of this service include the secure internal credential store for the logon of applications to remote systems (outbound connections), as well as all secure stores defined using the SAP HANA XS `$.security.Store` API. Every consumer of the service has its own system-internal application encryption key. Application encryption keys are encrypted with the root key of the data encryption service.

You should only change this root key if you need to ensure that it is not known outside your organization. Ideally, you change the root immediately after installation or receipt of your system from the hardware partner. At the latest, you must change it before any data is encrypted using the service. This means before you create any of the following things:

- A remote data source
- A HTTP destination
- An XS secure store
- A certificate collection with private key

#### Caution

It is important that you plan this root key change carefully as you will have to shut down the database. Not only that, but changing the root key after data has been encrypted will result in key information in the SSFS

and the database becoming inconsistent and encrypted data becoming inaccessible. Rectifying the problem could result in data loss. We recommend that you contact SAP Support if errors related to inconsistent SSFS or encryption failure occur.

Additionally, if the system supports multitenant database containers, you must change the root key before any tenant databases have been created.

## Procedure

1. Verify that no data has already been encrypted using the internal data encryption service by querying the following system views:

- CREDENTIALS (PUBLIC)
- P\_DPAPI\_KEY\_ (SYS)

### **i** Note

This view can only be accessed by user SYSTEM.

If the credential store is empty, then CREDENTIALS (PUBLIC) will also be empty. If there are no XS secure stores, then P\_DPAPI\_KEY\_ (SYS) will have no records with the caller XsEngine. If there are no certificate collections with private keys, then there will be no records with the caller PSEStore.

### **⚠** Caution

Do not proceed with the root key change if there is encrypted data.

2. Log on to the SAP HANA system host as the operating system user, <sid>adm.
3. Shut the system down using the sapcontrol program:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_no> -function Stop
```

4. Generate a new root encryption key using the hdbnsutil program.

```
cd /usr/sap/<sid>/HDB<instance_no>/exe  
./hdbnsutil - generateRootKeys --type=DPAPI
```

5. Start the system using the sapcontrol program:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_no> -function Start
```

6. Reset the consistency information in the SSFS using the hdbcons program:

```
cd /usr/sap/<sid>/HDB<instance_no>/exe  
./hdbcons "crypto ssfs resetConsistency"
```

The first time you execute the command, it does not reset the consistency information in the SSFS but outputs only a warning. To actually reset the consistency information in the SSFS, you must execute the command again within 20 seconds.

7. Change all application keys so that they are encrypted with the new root key.

You do this by executing the following SQL statement, for example using the SAP HANA studio or SAP HANA HDBSQL:

```
ALTER SYSTEM APPLICATION ENCRYPTION CREATE NEW KEY
```

#### **i** Note

You need `RESOURCE ADMIN` to execute this command.

After the next savepoint operation, new random internal application keys are created. New data is encrypted with the new application keys and the new keys are encrypted with the root encryption key. No re-encryption takes place. Any data encrypted with existing keys continues to be encrypted with these keys.

#### **i** Note

Depending on the workload of the database, the next savepoint operation may not happen for some time. You can force a savepoint with the statement: `ALTER SYSTEM SAVEPOINT`. In particular, you should force a savepoint with this statement if you plan to shut down your database right after generating the new application keys because a shutdown does not automatically write a savepoint and your change would be lost.

## Related Information

[Certificate Collections \[page 762\]](#)

[SAP HANA XS JavaScript API Reference](#)

### 6.4.1.5 Use FIPS 140-2 Certified Cryptographic Kernel in CommonCryptoLib

The SAP Cryptographic Library, CommonCryptoLib, supports a FIPS 140-2 compliant cryptographic kernel module. If required, you must manually install and enable it.

## 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. Install the FIPS 140-2 compliant crypto kernel as described in SAP Note 2117112.
2. Set the value of the parameter `[cryptography] ccl_fips_enabled` in the `global.ini` configuration file to **true**.
3. Restart the system or database.

## Results

The FIPS 140-2 certified crypto kernel, `libs1cryptokernel`, is used instead of the built-in crypto kernel, `libsapcrypto.so`.

If `libs1cryptokernel` 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

[Restart a System \[page 99\]](#)

[Stop and Start a Tenant Database \[page 134\]](#)

[SAP Note 2117112 - How to use the FIPS 140-2 certified Crypto Kernel with CommonCryptoLib](#)

[SAP Note 1848999 - Central Note for CommonCryptoLib 8 \(SAPCRYPTOLIB\)](#)

## 6.4.2 Client-Side Data Encryption (hdbuserstore)

The secure user store (`hdbuserstore`) is a tool installed with the SAP HANA client. You 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 only be used for SQLDBC and JDBC-based connections. The SAP HANA studio 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.

[Change the Secure User Store Encryption Key \[page 749\]](#)

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.

## 6.4.2.1 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.

## 6.5 Managing Encryption of Data Volumes in the SAP HANA Database

To protect data saved to disk from unauthorized access at operating system level, the SAP HANA database supports data encryption in the persistence layer.

[Enable Data Volume Encryption in an Existing SAP HANA System \[page 750\]](#)

There are two ways to enable data volume encryption in an existing SAP HANA system. The recommended way involves reinstalling your system. If this is not possible (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.

### [Change the Page Encryption Key \[page 755\]](#)

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.

### [Disable Data Volume Encryption \[page 756\]](#)

Disabling data volume encryption triggers the decryption of all encrypted data. Newly persisted data is not encrypted.

### [Data Volume Encryption in Multitenant Database Containers \[page 757\]](#)

Data volume encryption can be enabled individually for tenant databases in a multiple-container system.

## 6.5.1 Enable Data Volume Encryption in an Existing SAP HANA System

There are two ways to enable data volume encryption in an existing SAP HANA system. The recommended way involves reinstalling your system. If this is not possible (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.

### Context

To ensure that your database can always be restored to its most recent committed state, all data in the database is periodically copied to disk. You can ensure the privacy of data on disk by enabling data volume encryption. Ideally, you enable encryption immediately after installation or upgrade of SAP HANA. However, if you are already operating an SAP HANA database, it is also possible to enable data volume encryption and have your data encrypted retroactively.

Enabling data volume encryption does not increase data size.

Data volume encryption requires the availability of a cryptographic service provider on the SAP HANA server. The SAP Cryptographic Library (CommonCryptoLib) is available and used by default. For more information, see the *SAP HANA Security Guide*.

#### Caution

Do not enable data volume encryption if you plan to use the SAP HANA dynamic tiering option. It is not possible to create extended storage in encrypted SAP HANA databases. 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 and Options and Capabilities*.

### [Enable Data Volume Encryption with System Reinstallation \[page 751\]](#)

The recommended way to enable data volume encryption in an existing SAP HANA system is after reinstalling the system.

### [Enable Data Volume Encryption Without System Reinstallation \[page 753\]](#)

---

If it is not possible to reinstall your SAP HANA system to enable data volume 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.

## Related Information

[SAP Note 1848999 - Central Note for CommonCryptoLib 8 \(replacing SAPCRYPTOLIB\)](#)

[SAP Note 2093286 - Migration from OpenSSL to CommonCryptoLib \(SAPCrypto\)](#)

[Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities \[page 1360\]](#)

### 6.5.1.1 Enable Data Volume Encryption with System Reinstallation

The recommended way to enable data volume encryption in an existing SAP HANA system is after reinstalling the system.

#### Prerequisites

- You have the privileges required to perform an installation, as well as a backup and recovery.
- You have the system privilege RESOURCE ADMIN.
- If you are using the SAP HANA cockpit:
  - You have the privileges granted by the role `sap.hana.security.cockpit.roles::MaintainDataVolumeEncryption`.
  - The *Data Storage Security* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA Security Overview* tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.
- If the system is a secondary site in a system replication setup, verify how parameter replication is configured **before** attempting to enable data volume encryption. For more information about how parameter replication affects data volume encryption, see the SAP HANA Security Guide.

#### Context

Enabling data volume encryption after re-installing your system ensures that a new root encryption key is generated. In addition, it provides complete protection. If you enable encryption without a re-installation, due to the shadow memory nature of SAP HANA persistence, outdated versions of pages may still remain unencrypted on disk.

For more information about this recommendation, see SAP Note 2159014.

## Procedure

1. Perform a data backup.
2. Uninstall your system.

If possible, overwrite the former data area with random values.

3. Reinstall your system.

### ➔ Remember

If the system is a secondary site in a system replication setup, you must configure system replication now, before enabling data volume encryption.

4. Enable data volume encryption:

You can do this using either the SAP HANA cockpit or SAP HANA studio.

Option	Description
SAP HANA cockpit	<ol style="list-style-type: none"><li>1. Open the <i>Data Volume Encryption</i> app by clicking the <i>Data Storage Security</i> tile on the homepage of the SAP HANA cockpit.</li><li>2. Click the <i>Encrypt Data Volumes</i> button in the footer bar.</li></ol>
SAP HANA studio	<ol style="list-style-type: none"><li>1. In the Security editor of the system or database to be encrypted, choose the <i>Data Volume Encryption</i> tab.</li><li>2. Select <i>Activate encryption of data volumes</i> and choose  (<i>Deploy</i>).</li></ol>

5. Recover your system.

## Results

All data persisted to data volumes is encrypted. The status of data volume encryption is *Encrypted*.

### i Note

In the SAP HANA studio, you must refresh () the editor to see status changes.

For more information, see *SAP HANA SQL and System Views Reference* on SAP Help Portal.

**Task overview:** [Enable Data Volume Encryption in an Existing SAP HANA System \[page 750\]](#)

## Related Information

[Enable Data Volume Encryption Without System Reinstallation \[page 753\]](#)

[SAP Note 2159014 - FAQ: SQP HANA Security !\[\]\(30ff805cf73ae083425c51ec1695dcbc\_img.jpg\)](#)

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Creating Backups \[page 920\]](#)

[Recovering an SAP HANA Database \[page 949\]](#)

## 6.5.1.2 Enable Data Volume Encryption Without System Reinstallation

If it is not possible to reinstall your SAP HANA system to enable data volume 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 have changed the root encryption key if necessary. SAP HANA generates unique root keys on installation. However, if you received SAP HANA pre-installed from a hardware or hosting partner, you might want to change the root key used for data volume encryption to ensure it is not known outside your organization. For more information, see *Change the Root Encryption Key for Data Volume Encryption*.

#### **i** Note

In a system-replication configuration, change the root key used for data volume encryption in the primary system only. The new key will be propagated to all secondary systems.

- If you are using the SAP HANA cockpit:
  - You have the privileges granted by the role `sap.hana.security.cockpit.roles::MaintainDataVolumeEncryption`.
  - The *Data Storage Security* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA Security Overview* tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.
- If you are using the SAP HANA studio, you have the system privilege `RESOURCE ADMIN`.
- If the system is a secondary site in a system replication setup, verify how parameter replication is configured **before** attempting to enable data volume encryption. For more information about how parameter replication affects data volume encryption, see the SAP HANA Security Guide.

### Context

For maximum protection, we recommend that you reinstall your SAP HANA system before enabling data volume 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 the 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 will only be fully protected after some delay.

For more information about this recommendation, see SAP Note 2159014.

## Procedure

Enable data volume encryption:

You can do this using either the SAP HANA cockpit or SAP HANA studio.

Option	Description
SAP HANA cockpit	<ol style="list-style-type: none"><li>1. Open the <i>Data Volume Encryption</i> app by clicking the <i>Data Storage Security</i> tile on the homepage of the SAP HANA cockpit..</li><li>2. Click the <i>Encrypt Data Volumes</i> button in the footer bar.</li></ol>
SAP HANA studio	<ol style="list-style-type: none"><li>1. In the Security editor of the system or database to be encrypted, choose the <i>Data Volume Encryption</i> tab.</li><li>2. Select <i>Activate encryption of data volumes</i> and choose  (<i>Deploy</i>).</li></ol>

## Results

Encryption is now active for 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 encryption by service. Once encryption of a data volume has completed, the status changes to *Encrypted*.

### 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. This approach is therefore not recommended.

**Task overview:** [Enable Data Volume Encryption in an Existing SAP HANA System \[page 750\]](#)

## Related Information

[Enable Data Volume Encryption with System Reinstallation \[page 751\]](#)

[SAP Note 2159014 - FAQ: SQP HANA Security !\[\]\(deecc0eaa0ef374c6f18066e4fe103c3\_img.jpg\)](#)

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

---

[Change the Root Encryption Key for Data Volume Encryption \[page 744\]](#)

[Disable Data Volume Encryption \[page 756\]](#)

[Setting Up System Replication \[page 782\]](#)

## 6.5.2 Change the Page Encryption Key

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.

### 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.

For more information, see *SAP HANA SQL and System Views Reference* on SAP Help Portal.

## 6.5.3 Disable Data Volume Encryption

Disabling data volume encryption triggers the decryption of all encrypted data. Newly persisted data is not encrypted.

### Prerequisites

- You have the system privilege `RESOURCE ADMIN`.
- If you are using the SAP HANA cockpit:
  - You have the privileges granted by the role `sap.hana.security.cockpit.roles::MaintainDataVolumeEncryption`.
  - The *Data Storage Security* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA Security Overview* tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.
- If you are using the SAP HANA studio, you have the system privilege `RESOURCE ADMIN`.

### Procedure

Disable data volume encryption:

Option	Description
<b>SAP HANA cockpit</b>	<ol style="list-style-type: none"><li>1. Open the <i>Data Volume Encryption</i> app by clicking the <i>Data Storage Security</i> tile on the homepage of the SAP HANA cockpit.</li><li>2. Click the <i>Decrypt Data Volumes</i> button in the footer bar.</li></ol>
<b>SAP HANA studio</b>	<ol style="list-style-type: none"><li>1. In the Security editor of the system or database to be encrypted, choose the <i>Data Volume Encryption</i> tab.</li><li>2. Deselect the <i>Activate encryption of data volumes</i> checkbox and choose  (<i>Deploy</i>).</li></ol>

## Results

Encryption is now deactivated. Data starts being decrypted in the background. Depending on the size of the SAP HANA system, this process can be very time consuming. Only after this process has completed is all your data unencrypted. You can monitor the progress of decryption service by service. Once decryption of a data volume has completed, the status changes to *Unencrypted*.

### Note

In the SAP HANA studio, you must refresh () the editor to see status changes.

## Related Information

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

## 6.5.4 Data Volume Encryption in Multitenant Database Containers

Data volume encryption can be enabled individually for tenant databases in a multiple-container system.

Ideally, you enable encryption immediately after installation or upgrade of SAP HANA. This also applies to systems installed in multiple-container mode. Any subsequently created tenant databases will then automatically have encryption enabled. If a particular tenant database does not require encryption, the tenant database administrator can switch it off independently of the system in the Security editor of the SAP HANA studio or using the *Data Volume Encryption* app of the SAP HANA cockpit.

If encryption is not enabled after system installation, you can enable it retroactively either for all tenant databases together by making the setting in the system database, or for individual tenant databases by making the setting in the relevant tenant database.

### Caution

If you enable data volume encryption after a tenant database has been created and is already in operation, only the pages in use within the data volumes will be encrypted. Pages in the 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 will only be fully protected after some delay. To attain complete protection immediately, the overall process is:

1. Perform a data backup.
2. Drop the tenant database.
3. Clean the disk space.
4. Create the tenant database again.
5. Enable encryption.

---

6. Perform a data recovery.

## 6.6 Managing Client Certificates in the SAP HANA Database

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 systems that support multitenant database containers, in-database certificates are also used to secure communication during the process of copying or moving a tenant database between two systems. For more information, see *Copying and Moving Tenant Databases Between Systems* in the *SAP HANA Administration Guide*.

The following figure shows a typical certificate management workflow. A full separation of duties is possible through user authorization. For more information, see *SQL Statements and Authorization for In-Database Certificate Management*.



Certificate Management Workflow

You can manage certificates in the SAP HANA cockpit.

### **i** Note

Additional privileges are required to access the certificate management apps of the SAP HANA cockpit. These privileges are available in the roles delivered with the SAP HANA cockpit. The privileges for certificate management indicated above are partially included in these roles.

## Certificate Management in the File System

Although we recommend using in-database storage, it is possible to store and manage the certificates required for certificate-based user authentication and secure client-server communication in trust and key stores located in the file system.

### ➔ Recommendation

If you migrate from managing certificates in the file system to managing them in the database, delete all related files from the file system to avoid any potential conflicts. For more information, see SAP Note 2175664.

The certificates required to secure all internal communication channels and HTTP client access using SAP Web Dispatcher are contained in files located in the file system. In-database storage of certificates for these communication channels is not supported. Do not delete these files from the file system.

For more information about how to configure the usage of trust and key stores in the file system, see *Server-Side SSL Configuration Properties for External Communication* in the *SAP HANA Security Guide*.

## Overview of Certificate Handling

Certificates can be stored for...	...in the database	...in the file system
Secure client-server communication over JDBC/ODBC	Yes	Yes
Server client-server communication over HTTP	No	Yes
Secure internal communication	No	Yes
User authentication (SAML assertions, SAP logon and assertion tickets, X.509 certificates)	Yes	Yes

## Related Information

[SQL Statements and Authorization for In-Database Certificate Management \[page 770\]](#)

[SAP HANA Cockpit \[page 22\]](#)

[SAP Note 2175664 - Migration of file system-based X.509 certificate stores to in-database certificate stores](#)

## 6.6.1 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 data 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 \[page 762\]](#)

[Server-Side Data Encryption \[page 735\]](#)

### 6.6.1.1 Certificate Details

The *Certificate Store* app allows you to view the details of all certificates in the certificate store of the SAP HANA database.

Field	Description
<i>Issued To (CN)</i>	Common name of the person or entity identified by the certificate
<i>Issued To (DN)</i>	Distinguished name of the person or entity identified by the certificate
<i>Issued By (CN)</i>	Common name of the entity that verified the information and issued the certificate
<i>Issued By (DN)</i>	Distinguished name of the entity that verified the information and issued the certificate
<i>Issued On</i>	Date on which the certificate was issued
<i>Expires On</i>	End of certificate's validity
<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

Field	Description
<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 \[page 763\]](#)

## 6.6.2 Certificate Collections

A certificate collection (also referred to as a personal security environment 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.

### **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.

Certificate collections uniquely serve one of the following purposes in the database in which they exist:

- 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 the Secure Sockets Layer (SSL) protocol
- Database replication for multitenant database containers

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 SSL is being enforced for all client connections and the certificate collection used for SSL is deleted, no new client connections to the database can be opened.

## 6.6.2.1 Certificate Collection Details

The *Certificate Collections* app allows you to view the details of all certificate collections in the SAP HANA database.

Field	Description
<i>Purpose</i>	Purpose of the collection: <ul style="list-style-type: none"><li>• User authentication based on:<ul style="list-style-type: none"><li>◦ SAML assertions</li><li>◦ X.509 certificates</li><li>◦ Logon and assertion tickets</li><li>◦ JSON Web Token (JWT)</li></ul></li><li>• Client-server communication over JDBC/ODBC secured using the Secure Sockets Layer (SSL) protocol</li><li>• Database replication for multitenant database containers</li></ul>
<i>Provider</i>	SAML identity provider if the collection purpose is <i>SAML</i>
<i>Private Key</i>	Indicates whether or not a private key has been set for the collection  Only a collection with the purpose <i>SSL</i> requires a private key. This is the key that the SAP HANA server uses to identify itself to connecting clients.
<i>Created By</i>	Database user who created the collection
<i>Comment</i>	Optional comment

Field	Description
<a href="#">Certificates</a>	<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"> <li>• TRUST The certificate is the public-key certificate of a trusted communication partner.</li> <li>• PERSONAL The certificate is a server certificate belonging to the SAP HANA system and contains a private key.</li> <li>• 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.</li> </ul> <p>For more information about the other certificate fields, see <a href="#">Certificate Details</a>.</p>

## Related Information

[Certificate Details \[page 761\]](#)

### 6.6.3 View Certificates in the Certificate Store

You can view certificates stored in the database using the [Certificate Store](#) app of the SAP HANA cockpit.

#### Prerequisites

- You have the privileges granted by the role `sap.hana.security.cockpit.roles:DisplayCertificateStore`.
- The [Certificate Store](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the [SAP HANA Certificate Management](#) tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

#### Procedure

Open the [Certificate Store](#) app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.

---

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* app by clicking the collection name in the *Used In* column.

## Related Information

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Certificate Details \[page 761\]](#)

## 6.6.4 View Certificate Collections

You can view the certificate collections available in the database using the *Certificate Collections* app of the SAP HANA cockpit.

### Prerequisites

- You have the privileges granted by the role `sap.hana.security.cockpit.roles:DisplayCertificateStore`.
- The *Configure Certificate Collections* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA Certificate Management* tile catalog. For more information, see *Customizing the Homepage of SAP HANA Cockpit*.

### Procedure

Open the *Certificate Collections* app by clicking the *Configure Certificate Collections* tile on the homepage of the SAP HANA cockpit.

The *Certificate Collections* app 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

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Certificate Collections \[page 762\]](#)

[Certificate Collection Details \[page 763\]](#)

## 6.6.5 Import a Trusted Certificate into the Certificate Store

You store the public-key certificates of trusted communication partners, as well as the root certificates of trusted Certification Authorities directly in the SAP HANA database using the *Certificate Store* app.

### Prerequisites

- You have the privileges granted by the role `sap.hana.security.cockpit.roles:MaintainCertificates`.
- The certificate that you want to add is available on your client in PEM format.
- The *Certificate Store* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA Certificate Management* tile catalog. For more information, see *Customize the Homepage of SAP HANA Cockpit*.

### Procedure

1. Open the *Certificate Store* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.  
The *Certificate Store* app 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

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Certificate Details \[page 761\]](#)

## 6.6.6 Create a Certificate Collection

You create a certificate collection using the *Certificate Collections* app. Then, you add the relevant trusted certificates and if necessary, the server certificate.

### Prerequisites

- You have the privileges granted by the role `sap.hana.security.cockpit.roles:MaintainCertificateCollections`.
- 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.
- The *Configure Certificate Collections* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA Certificate Management* tile catalog. For more information, see *Customizing the Homepage of SAP HANA Cockpit*.

### Procedure

1. Open the *Certificate Collections* app by clicking the *Configure Certificate Collections* tile on the homepage of the SAP HANA cockpit.

The *Certificate Collections* app 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

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Server-Side Data Encryption \[page 735\]](#)

[Set the Purpose of a Certificate Collection \[page 768\]](#)

## 6.6.7 Set the Purpose of a Certificate Collection

You specify the purpose of a collection using the [Certificate Collections](#) app, 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 privileges granted by the role `sap.hana.security.cockpit.roles:EditCertificateStore`.
- You have the `REFERENCES` privilege on the certificate collection.

- To configure the collection for a user authentication purpose, you have the system privilege `USER ADMIN`.
- To configure the collection for the purpose `SSL` (secure client-server communication over JDBC/ODBC), you have the system privilege `SSL ADMIN`. In addition, the server certificate containing the server's private key must be part of the collection. For more information, see *Add a Server Certificate to a Certificate Collection*.
- To configure the collection for the purpose `DATABASE REPLICATION`, you have the system privilege `DATABASE ADMIN`. For more information, see *Copying and Moving Tenant Databases Between Systems* in the *SAP HANA Administration Guide*.
- The *Configure Certificate Collections* tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it from the *SAP HANA Certificate Management* tile catalog. For more information, see *Customizing the Homepage of SAP HANA Cockpit*.

## Procedure

1. Open the *Certificate Collections* app by clicking the *Configure Certificate Collections* tile on the homepage of the SAP HANA cockpit.
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
<i>SAML</i>	User authentication based on SAML assertions
<i>SAP LOGON</i>	User authentication based on logon and assertion tickets
<i>SSL</i>	Client-server communication over JDBC/ODBC secured using SSL/TLS
<i>X509</i>	User authentication based on X.509 client certificates
<i>DATABASE REPLICATION</i>	Communication between two multiple-containers systems via external SQL connections for the purposes of copying or moving a tenant database

### Note

Only those purposes that can be set and that you are authorized to set 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.

## Related Information

[Open SAP HANA Cockpit \[page 23\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

[Copying and Moving Tenant Databases Between Systems \[page 145\]](#)

## 6.6.8 SQL Statements and Authorization for In-Database Certificate Management

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> <b>i</b> Note You can also view certificates using the <i>Certificate Store</i> app of the SAP HANA cockpit.	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 which certificates are used in a certificate collection	<pre>SELECT * FROM PSE_CERTIFICATES</pre> <b>i</b> Note You can also see this information in the <i>Certificate Store</i> app of the SAP HANA cockpit.	Object privilege ALTER, DROP, or REFERENCES on the certificate collection
Add a certificate to the in-database certificate store	<pre>CREATE CERTIFICATE FROM &lt;certificate_content&gt; [ COMMENT &lt;comment&gt; ]</pre>	System privilege CERTIFICATE ADMIN
Delete a certificate from the in-database certificate	<pre>DROP CERTIFICATE &lt;certificate_id&gt;</pre> <b>i</b> Note If the certificate has already been added to a certificate collection, it can't be deleted.	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> <p><b>i Note</b> You can also view certificate collections using the <i>Certificate Collection</i> app of the SAP HANA cockpit.</p>	<p>System privilege CATALOG READ and either TRUST ADMIN, USER ADMIN, or SSL ADMIN</p> <p><b>i Note</b> 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>
Create a certificate collection	<pre>CREATE PSE &lt;PSE_name&gt;</pre>	System privilege TRUST ADMIN
Add a public-key certificate to a certificate collection	<pre>ALTER PSE &lt;PSE_name&gt; ADD CERTIFICATE &lt;certificate_id&gt;</pre>	<ul style="list-style-type: none"> <li>• Nothing if you're the owner of the certificate collection</li> <li>• Object privilege ALTER on the certificate collection if you're not the owner</li> </ul>
Remove a public-key certificate from a certificate collection	<pre>ALTER PSE &lt;PSE_name&gt; DROP CERTIFICATE &lt;certificate_id&gt;</pre> <p><b>i Note</b> 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>	
Add a private key to a certificate collection	<pre>ALTER PSE &lt;PSE_name&gt; SET OWN CERTIFICATE &lt;certificate_content&gt;</pre>	<ul style="list-style-type: none"> <li>• Nothing if you're the owner of the certificate collection</li> <li>• Object privilege ALTER on the certificate collection if you're not the owner</li> </ul>
Set the purpose of a certificate collection	<pre>SET PSE &lt;PSE_name&gt; PURPOSE &lt;PSE_purpose&gt;</pre> <p>The following PSE purposes are possible:</p> <ul style="list-style-type: none"> <li>• SAML</li> <li>• SAP LOGON</li> <li>• X509</li> <li>• SSL</li> <li>• JWT</li> </ul>	<ul style="list-style-type: none"> <li>• System privilege USER ADMIN or SSL ADMIN if you're the owner of the certificate collection USER ADMIN is needed if the purpose of the certificate collection is user authentication (SAML, X.509, or logon tickets), and SSL ADMIN is required if the purpose is secure client-server communication (SSL)</li> <li>• Object privilege REFERENCES on the certificate collection and system privilege USER ADMIN or SSL ADMIN if you're not the owner of the certificate collection</li> </ul> <p><b>i Note</b> If the purpose of the PSE is SSL, then it must already have a private key added.</p>

To...	Execute the Statement...	With the Authorization...
Unset the purpose of a certificate collection	<pre>UNSET PSE &lt;PSE_name&gt; PURPOSE &lt;PSE_purpose&gt;</pre>	<ul style="list-style-type: none"> <li>• System privilege SSL ADMIN if the purpose is secure client-server communication (SSL)</li> <li>• System privilege USER ADMIN for all other purposes</li> </ul>
Delete a certificate collection <div style="background-color: #fff9c4; padding: 5px; margin-top: 10px;"> <p><b>i Note</b></p> <p>If the certificate collection has already been assigned a purpose, it can't be deleted.</p> </div>	<pre>DROP PSE &lt;PSE_name&gt;</pre>	<ul style="list-style-type: none"> <li>• Nothing, if you're the owner of the certificate collection</li> <li>• Object privilege DROP on the certificate collection, if you're not the owner</li> </ul>

# 7 Availability and Scalability

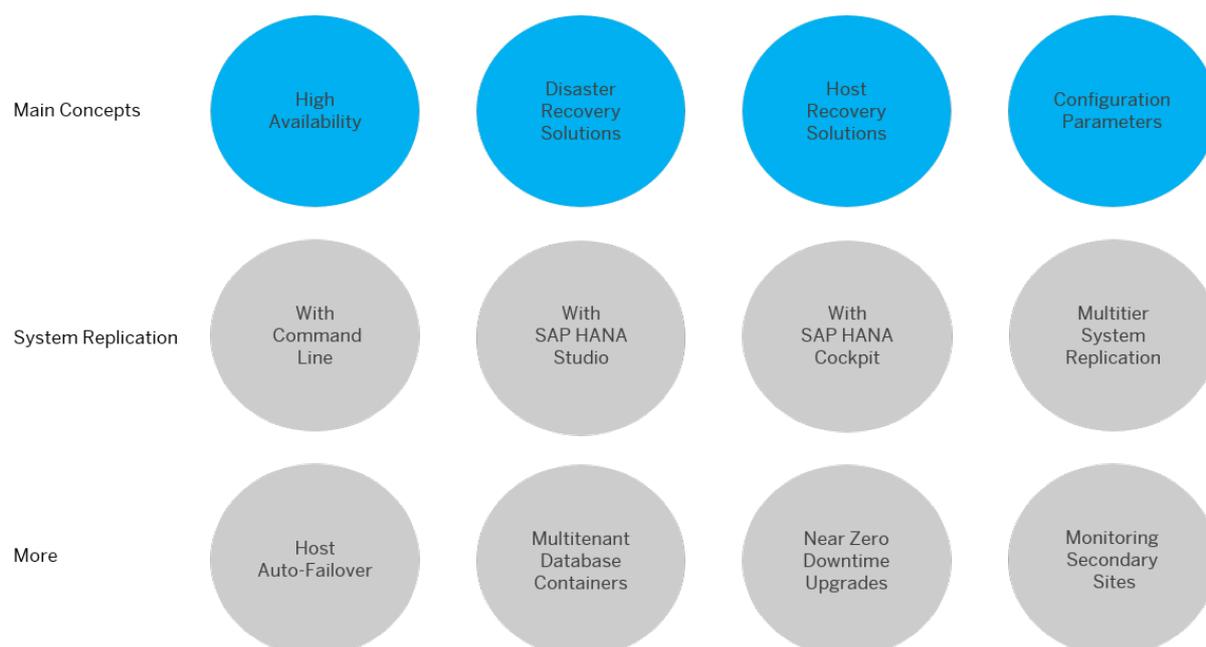
This section covers administration concepts and tasks for backup and recovery, scalability and high availability.

## 7.1 High Availability Navigation Support

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.

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.



- [High Availability for SAP HANA \[page 774\]](#)
- [SAP HANA Disaster Recovery Support \[page 776\]](#)
- [SAP HANA Fault Recovery Support \[page 779\]](#)
- [System Replication Configuration Parameters \[page 792\]](#)
- [Set up System Replication with hdbnsutil \[page 782\]](#)
- [Managing System Replication in the SAP HANA Studio \[page 828\]](#)
- [Monitor System Replication \[page 844\]](#)
- [Set up Multitier System Replication with hdbnsutil \[page 823\]](#)
- [Setting Up Host Auto-Failover \[page 850\]](#)

- [SAP HANA System Replication with Multitenant Database Containers \[page 803\]](#)
- [Use System Replication for Near Zero Downtime Upgrades \[page 816\]](#)
- [Monitoring Secondary Sites \[page 808\]](#)

## Related Information

[High Availability for SAP HANA \[page 774\]](#)

[SAP HANA Disaster Recovery Support \[page 776\]](#)

[SAP HANA Fault Recovery Support \[page 779\]](#)

[System Replication Configuration Parameters \[page 792\]](#)

[Set up System Replication with hdbnsutil \[page 782\]](#)

[Managing System Replication in the SAP HANA Studio \[page 828\]](#)

[Monitor System Replication \[page 844\]](#)

[Set up Multitier System Replication with hdbnsutil \[page 823\]](#)

[Setting Up Host Auto-Failover \[page 850\]](#)

[SAP HANA System Replication with Multitenant Database Containers \[page 803\]](#)

[Use System Replication for Near Zero Downtime Upgrades \[page 816\]](#)

[Monitoring Secondary Sites \[page 808\]](#)

## 7.2 High Availability for SAP HANA

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:

1. **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.
2. **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).

- 
3. 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.
  4. 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.

## 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.

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.

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 multitenant database containers on the system database level, this means the multitenant 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 multitenant database containers (including zero), also called tenant databases. For more information on multitenant database containers 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

[SAP Note 1999880](#)

[SAP Note 2183363](#)

[Creating and Configuring Tenant Databases \[page 104\]](#)

### 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](#)

## 7.2.1 SAP HANA Disaster Recovery Support

SAP HANA offers three levels of disaster recovery support - backups, storage replication and system replication.

### Backups

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.

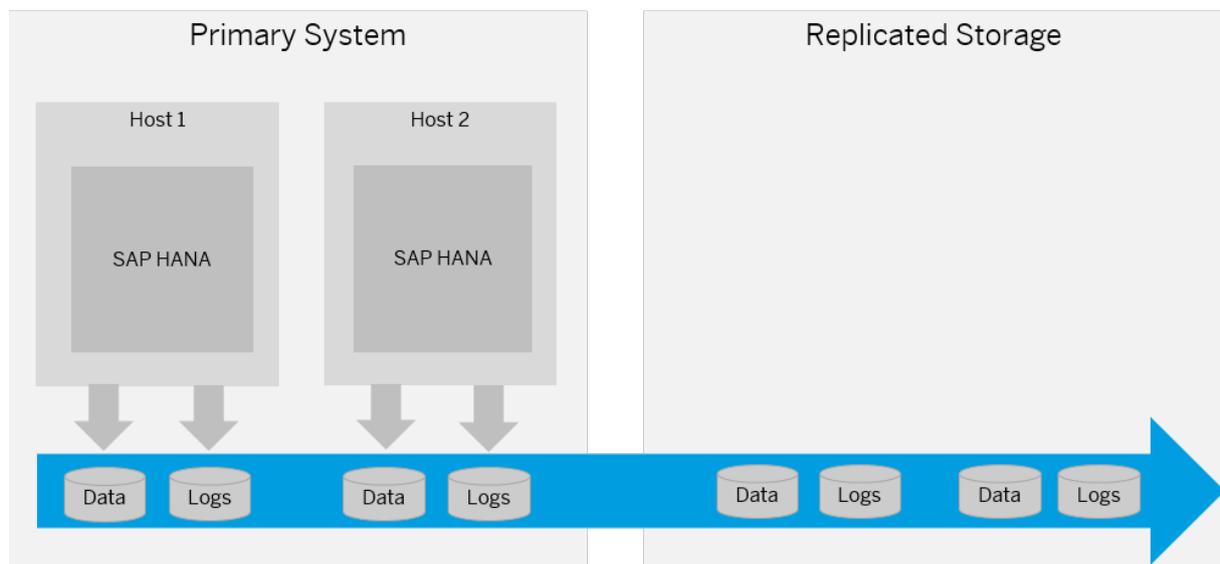
A savepoint is a periodic point in time, when all the changed data is written to storage, in the form of pages. One goal of performing savepoints is to speed up restart: when starting up the system, logs need not be processed from the beginning, but only from the last savepoint position. Savepoints are 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.

## Storage Replication

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.

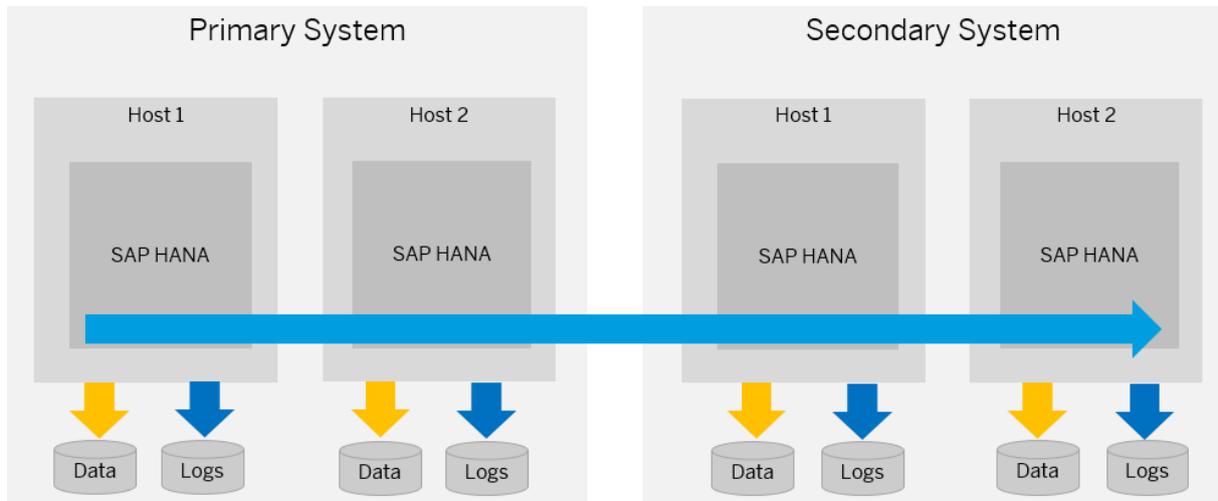


Due to its continuous nature, storage replication (sometimes also called remote storage mirroring) can be a more attractive option than backups, as it reduces the amount of time between the last backup and a failure. Another advantage of storage replication is that it also enables a much shorter recovery 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](#)

## System Replication

Usually system replication is set up so that a secondary standby 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 have one primary system and can have multiple secondary systems. Each service instance of the primary SAP HANA system communicates with a counterpart in the 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:

- Synchronous in-memory (default): The primary system commits the transaction after it receives a reply that the log was received by the secondary system, but before it has been persisted. The transaction delay in the primary system is shorter, because it only includes the data transmission time.
- Synchronous with full sync option means that log write is successful when the log buffer has been written to the logfile of the primary and the secondary instance. In addition, when the secondary system is disconnected (for example, because of network failure) the primary systems suspends transaction processing until the connection to the secondary system is re-established. No data loss occurs in this scenario
- Synchronous: The primary system does not commit a transaction until it receives confirmation that the log has been persisted in the secondary system. This mode guarantees immediate consistency between both systems, however, the transaction is delayed by the time it takes to transmit the data to and persist it in the secondary system.
- Asynchronous: 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 does not wait for confirmation from the secondary system.

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.

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

[SAP Note 1755396](#)

[Set up System Replication with hdbnsutil \[page 782\]](#)

[SAP HANA Database Backup and Recovery \[page 868\]](#)

[Recovery with System Replication \[page 970\]](#)

## 7.2.2 SAP HANA Fault Recovery Support

SAP HANA offer two levels of fault recovery support - service auto-restart and host auto-failover.

### Service Auto-Restart

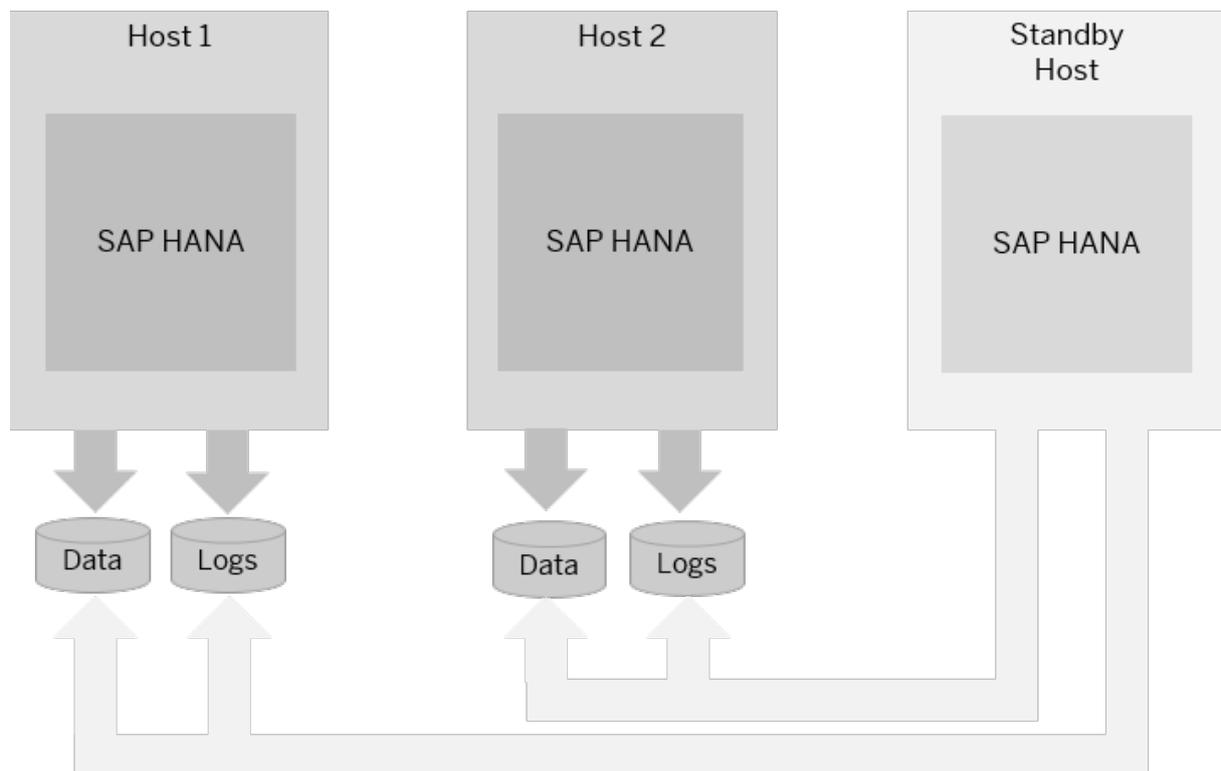
In the event of a software failure or an intentional intervention by an administrator that disables one of the configured SAP HANA services (Index Server, Name Server, and so on), the service will be restarted by the SAP HANA service auto-restart watchdog function, which automatically detects the failure 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.

## Host Auto-Failover

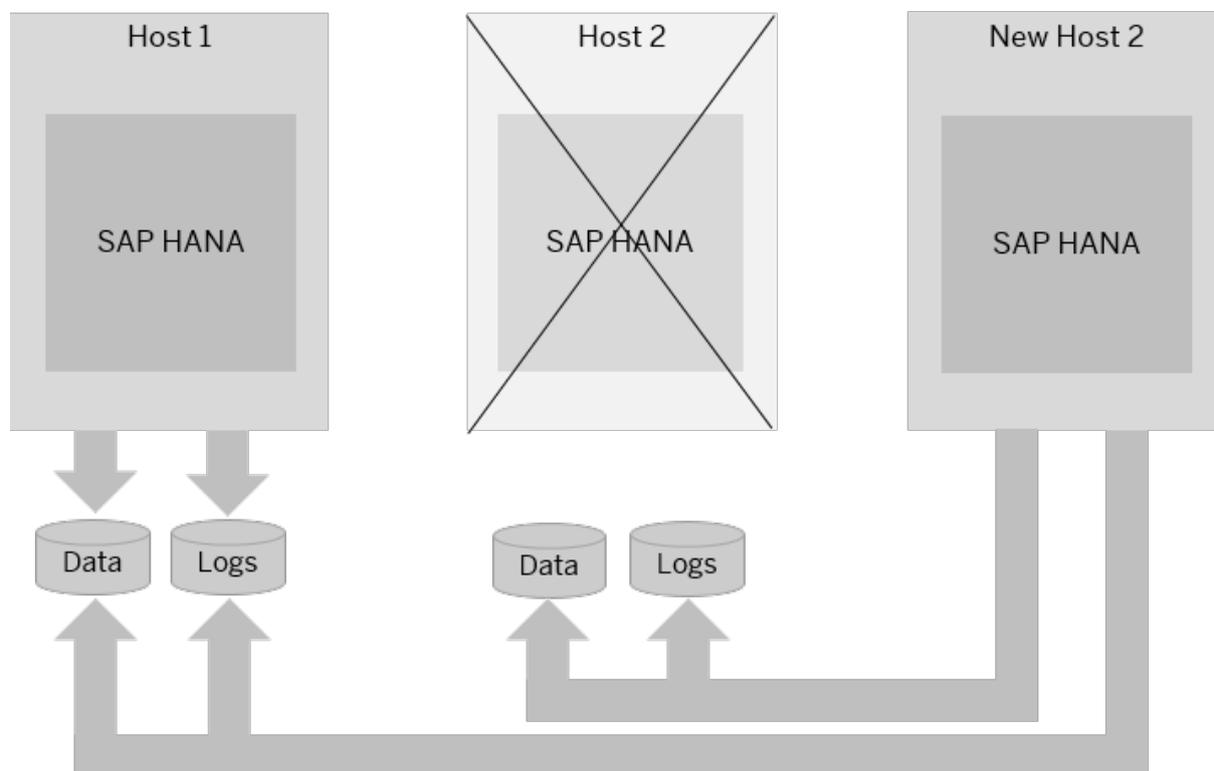
Host auto-failover is a local fault recovery solution that can be used in addition or as an alternative measure to system replication. One (or more) standby hosts are added to an SAP HANA system, and configured to work in standby mode. As long as they are in standby mode the databases on these hosts 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 a primary (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

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 \[page 850\]](#)

## 7.2.3 Setting Up System Replication

System replication can be set up or managed on the command line with `hdbnsutil`, using the SAP HANA studio, the SAP HANA cockpit or with SAP Landscape Virtualization Management (LVM).

There are a number of tools that can help you set up and manage system replication:

- `hdbnsutil`  
For more information see, *Set up System Replication with hdbnsutil*
- SAP HANA cockpit  
For more information see, *Managing System Replication in the SAP HANA Cockpit*
- SAP HANA studio  
For more information see, *Managing System Replication in the SAP HANA Studio*
- SAP Landscape Virtualization Management (LVM)  
For more information see, SAP Help Portal at <http://help.sap.com/nwlvms> > [Application Help](#) > [Managing System Landscapes](#) > [Performing Operations on Instances](#) >.

### **i** Note

There is no support for running one site with a little-endian system and the other site with a big-endian system in a system replication landscape. Some examples of endianness include:

- Intel: little-endian
- SLES 11 on Power: big-endian
- SLES 12 on Power: little-endian

## Related Information

[Set up System Replication with hdbnsutil \[page 782\]](#)

[Monitoring System Replication in the SAP HANA Cockpit \[page 843\]](#)

[Managing System Replication in the SAP HANA Studio \[page 828\]](#)

[SAP Landscape Virtualization Management \(LVM\) Documentation](#)

### 7.2.3.1 Set up System Replication with hdbnsutil

System replication enables recovery from a data center outage by switching to a secondary site. You can use the `hdbnsutil` tool to set up system replication.

## Prerequisites

- The primary and secondary system are both installed and configured. You have verified that both are independently up and running.

- The number of worker hosts in the secondary system has to be equal to the number of worker hosts in the primary system. 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.
- The software version of the secondary has to be equal or newer to the version on the primary.
- The secondary system must have the same SAP system ID, <SID> and `instance number` as the primary system.  
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 identical for both systems. Any changes made manually, or by SQL commands on one system must be manually duplicated on the other system.  
Automatic configuration parameter checks will alert you to configuration differences between the two systems.
- 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`  
See, *Rename an SAP HANA System Host*
- The required ports must be available. The same <instance number> is used for primary and secondary systems. The <instance number>+1 must be free on both systems, because this port range is used for system replication communications.
- 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.
- Data volume encryption must not be enabled in a secondary system before system replication is set up. Otherwise, the SSFS will become inconsistent and encrypted data inaccessible. If you want encryption on the secondary system, you can enable it after it's been integrated.  
For more information see, *Enable Data Volume Encryption in an Existing SAP HANA System* as well as *Encryption Key Management and Data Volume Encryption in the SAP HANA Security Guide*.

### **i** Note

If you do decide to enable data volume encryption after you configure system replication, it is important that you do not change the root key used for data volume encryption on any system involved.

## Context

As a general overview, the steps involved in setting up system replication between two systems, failing over to a secondary system, and failing back to a primary system are as follows:

## 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. Register the secondary system with the primary system (sr\_register).
  - e. 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.

## Related Information

[Rename an SAP HANA System Host \[page 585\]](#)

[Configuring the Network for Multiple Hosts \[page 1000\]](#)

[Implementing a HA/DR Provider \[page 854\]](#)

[Enable Data Volume Encryption in an Existing SAP HANA System \[page 750\]](#)

[Encryption Key Management \[page 739\]](#)

[SAP Note 2211663](#)

### 7.2.3.1.1 Configure the Primary System

You can configure the primary system for system replication with `hdbnsutil`.

#### Prerequisites

- You are logged on to the master name server's host as the `<sid>adm` user.
- If the hostnames of the primary and the secondary system are the same (for example, because two systems are used that have identical hostnames) change the hostnames used on the secondary system. See, *Rename an SAP HANA System Host*.

## Context

To set up system replication, the following configuration steps need to be carried out on the primary system.

You do this using the tool `hdbnsutil`, which initializes the topology of a database during the installation or exports, imports and converts the topology of an existing database, and the SAP HANA studio.

## Procedure

1. 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).
2. 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.
3. 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
--name	<primary_alias>	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
```

## Next Steps

Now you can configure the secondary system.

## Related Information

[Rename an SAP HANA System Host \[page 585\]](#)

---

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 922\]](#)

[Configure the Secondary System \[page 786\]](#)

## 7.2.3.1.2 Configure the Secondary System

To set up the secondary system for system replication you can configure and register it with the primary system using `hdbnsutil`.

### Prerequisites

- You have already configured a primary system, so that you can register a secondary system with it.
- You are logged on to the master name server's host as the root user on the secondary system.
- If the hostnames of the primary and the secondary system are the same (for example, because two systems are used that have identical hostnames) change the hostnames used on the secondary system. For more information see *Rename an SAP HANA System Host*.
- Check that hostname resolution works in both directions before configuring the secondary system for system replication. See SAP HANA Master Guide *Host Name Resolution for System Replication*.

### Context

There are different modes of log replication that can be used to send log information to the secondary instance. You need to decide which mode to use.

- Synchronous with full sync option (`mode=sync`. Full sync is configured with the parameter `[system_replication]/enable_full_sync`) means that log write is successful when the log buffer has been written to the log file of the primary and the secondary instance. In addition, when the secondary system is disconnected (for example, because of network failure) the primary systems suspends transaction processing until the connection to the secondary system is reestablished. No data loss occurs in this scenario.
- Synchronous (`mode=sync`) means the log write is considered as successful when the log entry has been written to the log volume of the primary and the secondary instance. When the connection to the secondary system is lost, the primary system continues 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.
- Synchronous in memory (`mode=syncmem`) means the log write is considered as successful, when the log entry has been written to the log volume of the primary and sending the log has been acknowledged by the secondary instance after copying to memory. When the connection to the secondary system is lost, the primary system continues transaction processing and writes the changes only to the local disk. 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 instance, but is more vulnerable to data loss.

- Asynchronous (mode=async): 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 does not wait for confirmation from the secondary system. 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.

Additionally, there are different operation modes for system replication.

- delta\_datashipping  
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.
- logreplay  
This mode does not require a delta data shipping anymore. Additionally the shipped redo log is continuously replayed on the secondary site.

To set up system replication, carry out the following configuration steps on the secondary system.

This is done using the tool `hdbnsutil`, which initializes the topology of a database during the installation or exports, imports and converts the topology of an existing database and using the SAP HANA studio.

## Procedure

1. 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).

2. Log on to the secondary system as user `<sid>adm` and use the `SAPControl` program to execute the following command to shut down the system:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StopSystem
HDB
```

3. 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
```

hdbnsutil Call Options

Option Name	Value	Description
--name	<secondary_alias>	Alias used to represent the secondary site

Option Name	Value	Description
--remoteHost	<primary_host>	Name of the primary host that the secondary registers with
--remoteInstance	<primary_instancnr>	Instance number of primary
--replicationMode	[sync syncmem async]	Log replication modes
--operationMode	[delta_datashipping logreplay]	Log operation mode

To check if the site has been successfully enabled for system replication with hdbnsutil run:

```
cd /usr/sap/<sid>/HDB<instancnr>/exe
```

```
./hdbnsutil -sr_state
```

4. 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
```

## Results

You have registered the secondary system with the primary system and system replication is enabled.

## Related Information

[Rename an SAP HANA System Host \[page 585\]](#)

[Monitoring System Replication \[page 804\]](#)

### 7.2.3.1.3 Enable Full Sync Option for System Replication

When activated the full sync option for 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.

## Related Information

[System Replication Command Line Reference \[page 800\]](#)

### 7.2.3.1.4 System Replication with Operation Mode Logreplay

System replication can be run in two operation modes: `delta_datashipping` or `logreplay`.

Since the first version of system replication, the operation mode `delta_datashipping` has been the default replication method. With the operation mode `logreplay` no delta data shippings are necessary anymore and the takeover time have been reduced and more components are already initialized at replication time..

Before you begin preparing a replication strategy for an SAP HANA system, you should be aware of the following important points regarding the operation mode `logreplay`.

- 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.
- 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`.

- Switching operation modes from logreplay back to delta\_datashipping requires a full data shipping.
- The combination of logreplay and delta\_datashipping in an SAP HANA multitier system configured with system replication is not supported.
- When the connection to the secondary with operation mode logreplay is not available, the primary system will keep 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.
- To prevent the log volume from running full the following criteria apply:
  - 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).
- The operation mode logreplay does not support history tables.

## Related Information

[System Replication Command Line Reference \[page 800\]](#)

### 7.2.3.1.4.1 Log Retention

With the operation mode `logreplay`, 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` 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` 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 `[system_replication]/logshipping_max_retention_size` 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).

## 7.2.3.1.5 System Replication Configuration Parameters

Several configuration parameters are available for configuring 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.

### **i** 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	120	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"> <li>• A snapshot has to exist on the new secondary when it starts up for the first time as secondary</li> <li>• The snapshot has to be compatible with the persistence of the new primary.</li> </ul> <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	Number of seconds, the primary waits for the acknowledge after sending a log buffer to the secondary site. If the primary does not receive the acknowledge 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

Parameter	Type	Unit	Default	System	Description
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>
logshipping_async_wait_on_buffer_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>

Parameter	Type	Unit	Default	System	Description
enable_full_sync	bool	bool	false	Primary	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.
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>

Parameter	Type	Unit	Default	System	Description
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>
keep_old_style_alert	bool	true/false	true	Primary	<p>Before SPS 09 closed replication connections and configuration parameter mismatches were alerted with statistics server 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.</p>

Parameter	Type	Unit	Default	System	Description
operation_mode	enum		delta_data_shipping	Secondary	<p>Operation mode of the secondary site during replication. There are two different settings for this parameter:</p> <ul style="list-style-type: none"> <li> <b>delta_data_shipping</b>  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. </li> <li> <b>logreplay</b>  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. . </li> </ul>

Parameter	Type	Unit	Default	System	Description
enable_log_retention	enum		auto	Primary	<p>Enables/Disables log retention on a system replication primary 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). There are three configuration options:</p> <ul style="list-style-type: none"> <li>• 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.</li> <li>• on Log retention is enabled</li> <li>• off Log retention is disabled</li> </ul> <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>

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 on primary side for syncing a system replication secondary system. This value only has an effect, if log retention is enabled. 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>

## Related Information

[Change a System Property \[page 217\]](#)

### 7.2.3.1.6 System Replication Command Line Reference

System replication commands and options.

sr\_commands

Command	Options	System	Online/Offline	Description
-sr_enable	[--name=<site alias>]	Primary	Online	Enables a site to serve as a system replication source site.  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.
-sr_disable		Primary	Online	Disables system replication capabilities on source site.
-sr_register	--remoteHost=<primary master host>	Secondary	Offline	Registers a site to a source site and creates the replication path for the system replication.
	--remoteInstance=<primary instance id>			
	--replicationMode=sync syncmem async			Specifies the replication mode.
	--operationMode=delta_datashipping logreplay			Specifies the operation mode.
	--name=<unique site name>			Specify the site name.
	[--force_full_replica]			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_initialize	--database=<database name> --volume=<volume id>			Initializes a given database or specific volume for system replication.
	[--force_full_replica]			If parameter is given, a full data shipping is initiated. Otherwise a delta data shipping is attempted.
-sr_changemode	--mode=sync syn-cmem 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](#) 

## 7.2.3.1.7 SAP HANA System Replication with Multitenant Database Containers

The usual SAP HANA system replication principles apply for multitenant database containers.

Before you begin preparing a replication strategy for an SAP HANA multiple-container system, you should be aware of the following important points.

- SAP HANA multiple-container 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 a tenant database is recovered on the primary system, the replication will be in an inconsistent state with the secondary site until this tenant database is reinitialized (for more information see `-sr_initialize` in the section *System Replication Command Line Reference* of the SAP HANA Administration).
- 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 multiple-container 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 multiple-container 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 multiple-container 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 multitenant database containers. 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 multiple-container systems using system replication technology. 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*.

- For SAP HANA multiple-container 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 \[page 773\]](#)

[SAP HANA Database Backup and Recovery \[page 868\]](#)

[System Replication Command Line Reference \[page 800\]](#)

[Copying and Moving Tenant Databases Between Systems \[page 145\]](#)

[Increase the System Isolation Level \[page 105\]](#)

### 7.2.3.1.8 Monitoring 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 in the Administration editor of the primary system as follows:

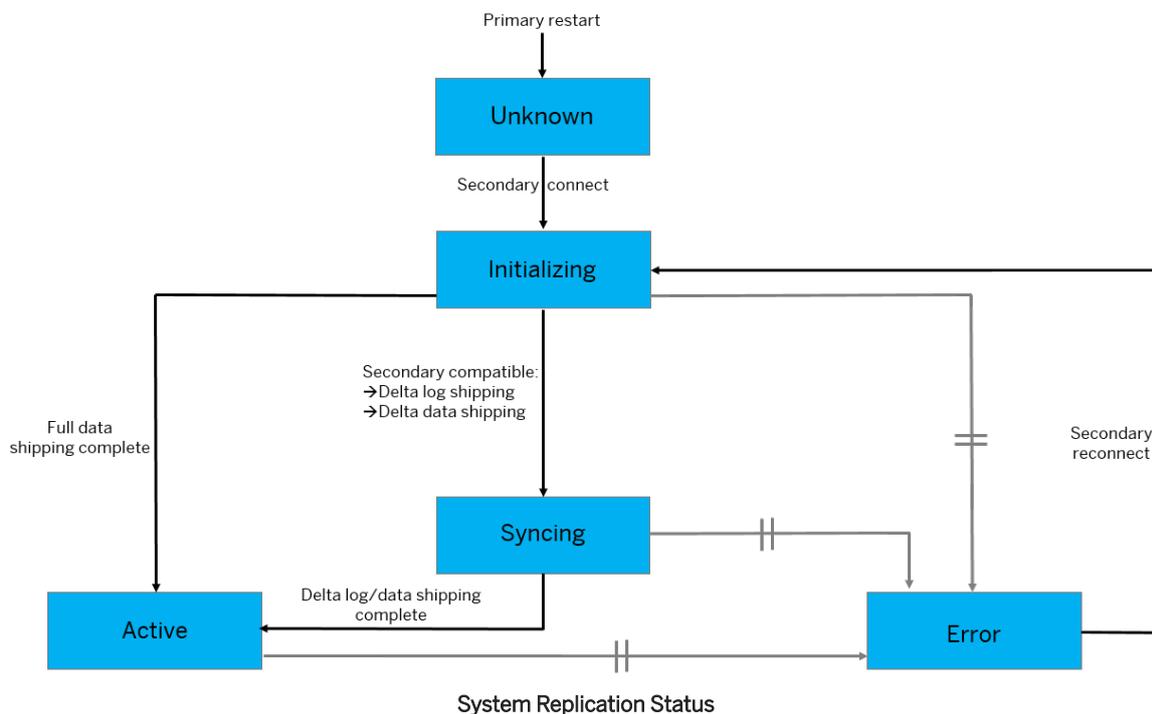
- The general status is displayed on the *Overview* tab.
- Detailed information is available on the **► Landscape ► System Replication ►** tab. Since the secondary instance does not accept SQL connections while data replication is active, basic information about the secondary system is also shown. For more information about the meaning of the individual fields, see the system view M\_SERVICE\_REPLICATION.

## Replication Status

The current status of system replication is shown in the field M\_SERVICE\_REPLICATION.SYSTEM\_STATUS

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.

Status	Description
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 and 79 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

## Related Information

[Monitoring Status and Resource Usage of System Components \[page 234\]](#)

[Monitoring Alerts \[page 242\]](#)

[Configure E-Mail Notifications for Alerts \[page 249\]](#)

[Monitoring INI File Parameter Changes \[page 806\]](#)

[System Replication Alerts \[page 807\]](#)

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

## 7.2.3.1.8.1 Monitoring 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.

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.

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_[infile name|*][/<LAYER>] = [section with  
wildcards|*][parameter with wildcards|*], ...  
<LAYER> := SYSTEM\|HOST\|DATABASE\|\"
```

## Related Information

[Configuring SAP HANA System Properties \(INI Files\) \[page 212\]](#)

### 7.2.3.1.8.2 System Replication Alerts

Alerts issued by the primary system warn you of potential problems.

The following alerts are issued by the primary system:

- Alert ID 78: System Replication Connection Closed
- Alert ID 79: System Replication Configuration Parameter Mismatch
- Alert ID 94: System Replication Logreplay Backlog

Alerts 78 and 79 are raised when a system replication connection is closed or when there is a system replication configuration parameter mismatch.

Starting with 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).

Before SPS09 there was one alert, categorized as an "Internal Event" (Alert 21). It was created when:

- The connection to the secondary site was closed.
- There was a configuration parameter mismatch between the replication sites.

Both situations were covered by one event type and could only be distinguished by the information text provided.

Since SP11 old style alerts based on alert 21 are not created anymore as a default.

You can create them by setting the configuration parameter `keep_old_style_alert` to `<true>` in the system replication section of the `global.ini` 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.

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.

The alert has a different priority based on the threshold reached:

- 10 GB: low
- 50 GB: medium
- 500 GB: high

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.

For information on alerts issued by hosts of the secondary system, see *Monitoring Secondary Sites*.

## Related Information

[SAP Note 1917938](#)

[Monitoring Secondary Sites \[page 808\]](#)

### 7.2.3.1.8.3 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.

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 \[page 302\]](#)

## 7.2.3.1.9 Takeover Decision

In addition to tools that may be used to monitor the overall system status when system replication is enabled a script is provided with SAP HANA that helps you decide when a takeover should be carried out.

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, a script called `landscapeHostConfiguration.py` is provided so that SAP HANA itself can communicate 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.

### i 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.

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@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 | no | warning | failover to ld8522 | | | 2 | default | default | master 2 | slave | worker | slave |
| ld8522 | 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@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 | no | info | | | | 0 | default | default | master 2 | slave | worker | standby |
| ld8522 | 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@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 | no | error | | | | 2 | default | default | master 2 | slave | worker | slave |
| ld8522 | no | ignore | | | | 0 | default | default | master 3 | slave | standby | standby |
overall host status: error
```

The return codes of the script are:

Return Codes

Return code	Description
0	Fatal Internal script error, the state could not be determined
1	Error
2	Warning
3	Info
4	OK

### **i** 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.

To help you decide if a takeover is advisable see the decision tree in [SAP Note 2063657](#)

## Related Information

[SAP Note 2063657](#) 

### 7.2.3.1.10 Perform a Takeover

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.

## 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, meaning it now accepts and responds to queries, using the following steps:

## 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 \[page 98\]](#)

[Monitoring SAP HANA Systems During Stop and Start \[page 103\]](#)

### 7.2.3.1.11 Performing a Failback

After a takeover has been carried out, and the data center is back in operation the roles between primary and secondary can be switched over. In this case, the former primary now has to be registered as the secondary with the active primary system.

This is the same procedure as is used for setting up a normal secondary described in *Configuring the Secondary System*. However, in this scenario when the new secondary is registered with the new primary it checks if it is possible to do a delta shipping rather than carrying out a full data shipping. If this is possible it only ships the delta which significantly reduces the initialization time.

When the new secondary starts up it checks first if there is a local snapshot available from 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 check are positive the new secondary can be initialized with a delta replica from the new primary.

---

## Related Information

[Configure the Secondary System \[page 786\]](#)

### 7.2.3.1.12 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 persistence 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 shipping.

After a new registration of the secondary site, delta data shipping is always attempted. To do this, it has to be ensured, that the persistence of the primary and secondary system are compatible. Compatibility is checked for current persistence as well as replication snapshots. They are checked for compatibility in the order listed below:

1. Current savepoint (last written savepoint):
  - This is available for every persistence
  - Can be used if the secondary has just been shutdown for a while
2. The active replication snapshot:
  - This is written on the system replication primary sites after successful data shippings.
  - Can be used if an old primary was used to prepare a failback.
3. The most recent replication snapshot:
  - This is written on the system replication primary and secondary sites for data shipping
  - It can be used, for example, when a secondary site has been taken over and is being re-registered again afterward.

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.

---

## 7.2.3.1.12.1 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.

## 7.2.3.1.13 Example Set Up of 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 \[page 585\]](#)

[SAP Note 611361](#)

### 7.2.3.1.14 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.

#### Procedure

1. Unregister the secondary system 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](#)

## 7.2.3.1.15 Use System Replication for Near Zero Downtime Upgrades

You can use 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 further reduce the downtime with the optimized update procedure. For more information, see *Prepare an Update for Reduced SAP HANA System Downtime*.

The secondary system can be initially installed with the new software version or upgraded to the new software version when 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.

### 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 (<sqlport>). Execute this command on the primary and secondary systems:

```
hdbuserstore SET SRTAKEOVER <publichostname>:<sqlport> <myrepointer>
<myrepointer_password>
```

Create a <myrepointer> 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:30015 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.

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.

From your installation directory execute as root:

```
./hdblcm --action=update
```

3. 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.

4. 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)

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. Stop the primary system
7. Upgrade the original primary from the installation directory as root user using the option --hdbupd\_server\_nostart. 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

If you are upgrading to revision 93 or greater see *SAP Note 2164506* for more information.

8. Use the SAP HANA lifecycle management tools to upgrade all the other components to the same revision as the server software.

9. 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]
```

10. Start the original primary.

## Related Information

[SAP Note 1984882](#)

[SAP Note 2164506](#)

[SAP Note 2386973](#)

### 7.2.3.1.16 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 \[page 541\]](#)

[Remove Hosts Using the Command-Line Interface \[page 549\]](#)

### 7.2.3.1.17 Add a New Host Manually to a Replicated System

You can add manually a new host to a replicated system.

#### Context

To add a new host manually to a replicated system perform the following steps (as `<sid>adm` on the command line):

#### Procedure

1. Stop the secondary system:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StopSystem  
HDB
```

2. Unregister the secondary system:

```
cd /usr/sap/<sid>/HDB<instancenr>/exe
```

```
./hdbnsutil -sr_unregister
```

3. Add the new host to the primary system (the primary system is still running):

For more information, see *Add Hosts Using the Command-Line Interface*

4. Start the secondary system:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StartSystem  
HDB
```

5. Add a new host to secondary system.

For more information, see *Add Hosts Using the Command-Line Interface*

6. Stop the secondary system:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StopSystem  
HDB
```

7. Re-register the secondary system:

```
cd /usr/sap/<sid>/HDB<instancenr>/exe
```

```
./hdbnsutil -sr_register ...
```

8. Start up the secondary system again:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StartSystem  
HDB
```

## Related Information

[Add Hosts Using the Command-Line Interface \[page 541\]](#)

[Remove Hosts Using the Command-Line Interface \[page 549\]](#)

### 7.2.3.1.18 Client Connection Recovery

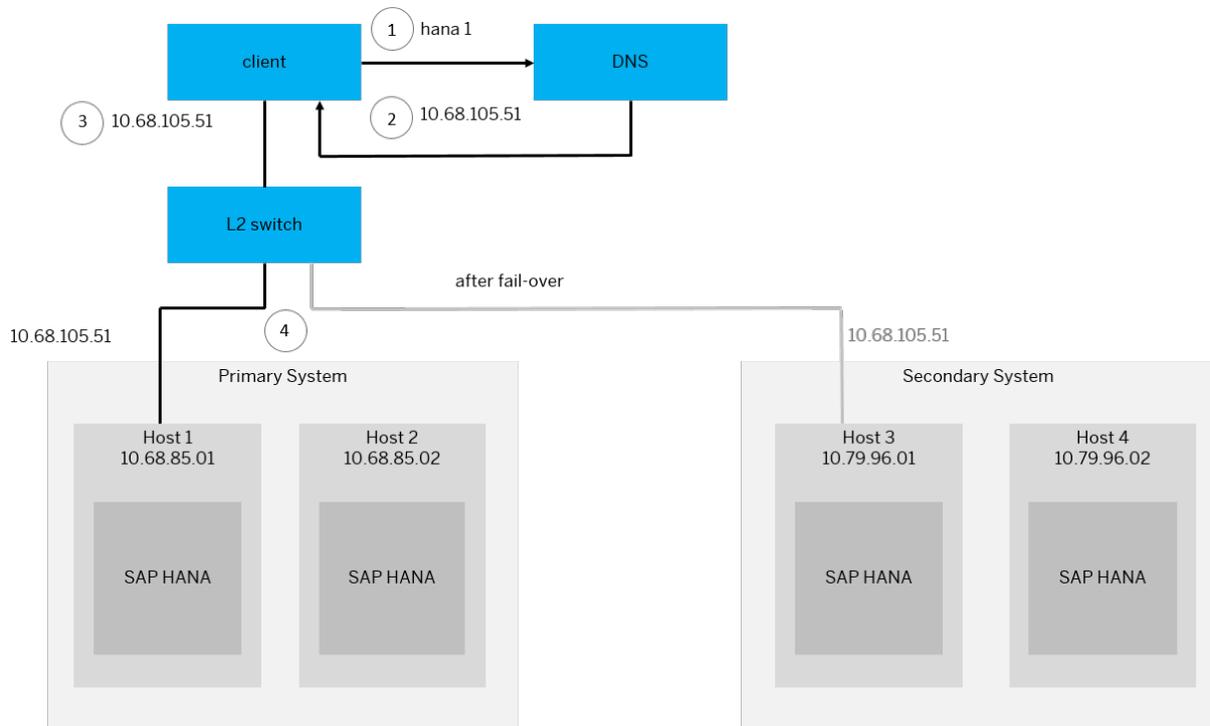
To allow for continued client communication with the SAP HANA system your high availability solution has to also support client connection recovery. Connection recovery after disaster recovery can be done with network-based IP redirection or network-based DNS redirection.

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 (10.68.104.51), 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 fail-over procedure, a script is executed which re-maps the unchanged logical IP address to the corresponding fail-over 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:



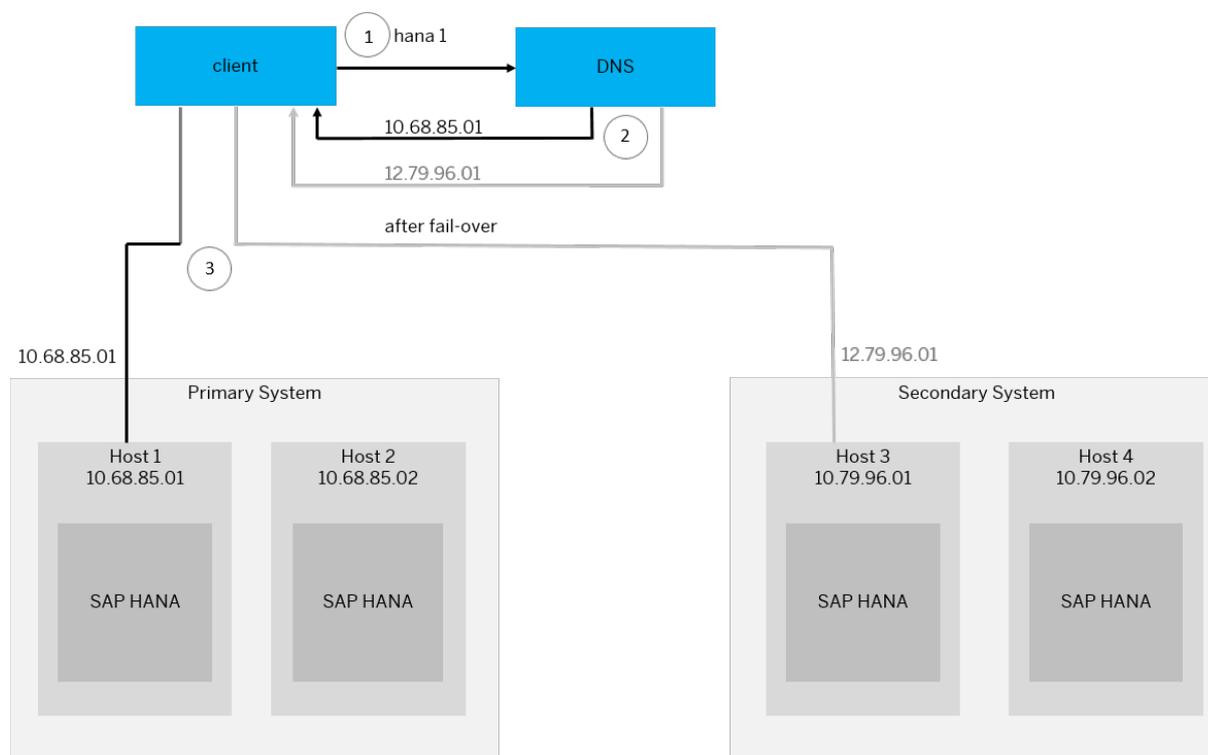
IP redirection can be implemented using a number of actual techniques, for instance with the use of Linux commands which affect the network ARP tables (`ip addr add/del...`), by configuring L2 network switches directly, or by using cluster management software. Following the IP redirection configuration, 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 fail-over 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 are currently working on or already 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 \[page 854\]](#)

[SAP Note 2053504](#)

### 7.2.3.1.19 Set up Multitier System Replication with hdbnsutil

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.

#### Prerequisites

- All systems must be installed with the same SID and have the same instance number as the primary system.

#### Context

To offer higher levels of availability you can link together multiple systems in a multitier system replication landscape. After setting up a basic system replication scenario you now 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 supports various replication mode combinations. For more information, see *Supported Replication Modes between Sites*.

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

[Supported Replication Modes between Sites \[page 840\]](#)

## 7.2.3.19.1 Takeover and Failback in 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.

### 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

1. SiteA failed, SiteB has taken over and now you attach SiteA as the tier 3 secondary.

Multitier system replication supports various replication mode combinations. For more information, see *Supported Replication Modes between Sites*.

- a. [C] Change the replication mode of the new tier 2 secondary:

```
cd /usr/sap/<sid>/HDB<instance_number>/exe
./hdbnsutil -sr_changemode --replicationMode=sync
```

- b. [C] Enable SiteC as the replication source:

```
hdbnsutil -sr_enable
```

- c. [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
```

- d. [A] Register SiteA as a new tier 3 secondary.

```
hdbnsutil -sr_register --replicationMode=async --name=SiteA --
remoteInstance=<instId> --remoteHost=<hostname_of_C>
```

- e. [A] Start the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function
StartSystem HDB
```

- f. [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.

2. You want to restore the original multitier setup

- a. [C] Stop the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function  
StopSystem HDB
```

- b. [C] Unregister SiteC from SiteB:

```
hdbnsutil -sr_unregister
```

- c. [A] Stop the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function  
StopSystem HDB
```

- d. [A] Register as secondary:

```
hdbnsutil -sr_register --replicationMode=sync --name=SiteA --  
remoteInstance=<instId> --remoteHost=<hostname_of_B>
```

- e. [A] Start the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function  
StartSystem HDB
```

- f. [B] Check in M\_SERVICE\_REPLICATION that sync system replication is ACTIVE from SiteB to SiteA.

- g. [A] SiteA takes over as the primary system:

```
hdbnsutil -sr_takeover
```

- h. [B] Stop the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function  
StopSystem HDB
```

- i. [A] Enable system replication:

```
hdbnsutil -sr_enable --name=SiteA
```

- j. [B] Register SiteB as the tier 2 secondary of SiteA.

```
hdbnsutil -sr_register --replicationMode=sync --name=SiteB --  
remoteInstance=<instId> --remoteHost=<hostname_of_A>
```

- k. [B] Start the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function  
StartSystem HDB
```

- l. [B] Enable SiteB as a replication source system:

```
hdbnsutil -sr_enable
```

- m. [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>
```

- n. [C] Start the SAP HANA database

```
/usr/sap/hostctrl/exe/sapcontrol -no <instance_number> -function  
StartSystem HDB
```

- 
- o. [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.

## Related Information

[Supported Replication Modes between Sites \[page 840\]](#)

### 7.2.3.1.20 Configuring Hostname Resolution for SAP HANA System Replication

To enable communication over a particular network you can define a mapping from an IP address to a hostname for SAP HANA system replication, which are used exclusively for SAP HANA system replication. This lets you use virtual hostnames on both systems without having to adapt the `/etc/hosts` file or where different internal and external hostnames need to be used.

You map IP addresses to hostnames by editing the section `system_replication_hostname_resolution` in the `global.ini` file on the *Configuration* tab of the SAP HANA studio. In the case of multitier system replication only the direct neighbors have to be specified in the mapping.

This technique can also be used if the hostnames have a domain suffix, for example. the internal hostnames `ab820*` and `ab830*`, but the public names have to include the domain, for example, `ab820*.abc.xyz.com` and `ab830*.def.xyz.com`

For more information, see *Host Name Resolution for System Replication*.

## Related Information

[Configure the Secondary System \[page 786\]](#)

### 7.2.3.1.21 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.

lz4 has been selected as the compression algorithm, because of its speed and compression ratios. Therefore 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

[LZ4](#) 

### 7.2.3.2 Managing System Replication in the SAP HANA Studio

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.

You can manage system replication in the SAP HANA studio. The following administration activities are possible:

- 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 failover to the secondary system in the event of a disaster and failback once the primary system is available again.
- Disabling system replication.

## Related Information

[Set Up System Replication \[page 829\]](#)

[Fail Over to the Secondary System \[page 833\]](#)

[Fail Back to the Original Primary System \[page 834\]](#)

[Disable System Replication \[page 837\]](#)

### 7.2.3.2.1 Set Up System Replication

To set up 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 installed and configured two identical, independently-operational SAP HANA systems – a primary system and a secondary system.

The secondary system must meet the following criteria with respect to the primary system:

- It must have a different host name, or host names in the case of a distributed system.  
If the host names of the primary and the secondary systems are the same (for example, because two systems that have identical host names are used) in the SAP HANA lifecycle management tool `hdblicm(gui)` to change the host name(s) of the secondary system.
- It must have the same number of worker hosts. This implies that if there is a standby host on the primary system it need not be available on the secondary system.
- It must have the same software version or higher.
- It must have the same SAP system ID (SID).

#### **i** Note

The primary replicates all relevant license information to the secondary. An additional license is not required.

- It must have the same system configuration in the system properties files (\*.ini files).

#### **i** Note

Any changes made manually or by SQL on one system must be manually duplicated on the other system. Automatic configuration parameter checks will alert you to configuration differences between the two systems.

- It should run on the same endianness platform.
- The required ports must be available.  
The same instance number is used for primary and secondary systems. The instance number +1 must be free on both systems, because this port range is used for system replication communications.
- 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 added both 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 both systems.  
You can do this in the Administration editor (*Configuration* tab) of the SAP HANA studio.
- 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.
- You have stopped the secondary system.
- Data volume encryption must not be enabled in a secondary system before system replication is set up. Otherwise, the SSFS will become inconsistent and encrypted data inaccessible. If you want encryption on the secondary system, you can enable it after it's been integrated.  
For more information see, *Enable Data Volume Encryption in an Existing SAP HANA System* as well as *Encryption Key Management* and *Data Volume Encryption* in the *SAP HANA Security Guide*.

### **i** Note

If you do decide to enable data volume encryption after you configure system replication, it is important that you do not change the root key used for data volume encryption on any system involved.

## 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*.
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:

Log Replication Mode	Description
Synchronous in-memory (default)	<p>If you select this option, the primary system commits the transaction after it receives confirmation that the log has been received by the secondary system but before it has been persisted. The transaction delay in the primary system is shorter because it only includes the data transmission time.</p> <p>This option provide 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 if both systems fail at the same time.</p>
Synchronous with full sync option	<p>Synchronous with full sync option means that log write is successful when the log buffer has been written to the logfile of the primary and the secondary instance. In addition, when the secondary system is disconnected (for example, because of network failure) the primary systems suspends transaction processing until the connection to the secondary system is re-established. No data loss occurs in this scenario.</p>
Synchronous	<p>If you select this option, the primary system does not commit a transaction until it receives confirmation that the log has been persisted in the secondary system.</p> <p>This option guarantees immediate consistency between both systems; no loss of data is guaranteed. However, the transaction is delayed by the time it takes to transmit the data to and persist it in the secondary system.</p>
Asynchronous	<p>If you select this option, 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 does not 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>

- f. 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 system appears as operational (■). The secondary system appears as operational (■) but with an error (✖) indicating that no connection to the database is available.

## Related Information

[SAP Note 611361](#)

[Add an SAP HANA System \[page 70\]](#)

[Stop a System \[page 98\]](#)

[System Replication Configuration Parameters \[page 792\]](#)

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 922\]](#)

[Rename an SAP HANA System Host \[page 585\]](#)

[Enable Data Volume Encryption in an Existing SAP HANA System \[page 750\]](#)

[Encryption Key Management \[page 739\]](#)

## 7.2.3.2.2 Fail Over to the Secondary System

If your primary data center is not available, due to a disaster or for planned downtime for example, you can fail over to a secondary data center by performing a takeover on your secondary system.

### 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.

## 7.2.3.2.3 Fail Back to the Original Primary System

Once your original primary data center is operational again after a disaster or planned downtime, for example, you can fail back to your original primary system.

### 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 two systems back to their original configuration. To do so, the original primary system will have to be started as secondary system and after both systems are back in sync, you 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:

Log Replication Mode	Description
Synchronous in-memory (default)	<p>If you select this option, the primary system commits the transaction after it receives confirmation that the log has been received by the secondary system but before it has been persisted. The transaction delay in the primary system is shorter because it only includes the data transmission time.</p> <p>This option provide 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 if both systems fail at the same time.</p>
Synchronous	<p>If you select this option, the primary system does not commit a transaction until it receives confirmation that the log has been persisted in the secondary system.</p> <p>This option guarantees immediate consistency between both systems; no loss of data is guaranteed. However, the transaction is delayed by the time it takes to transmit the data to and persist it in the secondary system.</p>
Asynchronous	<p>If you select this option, 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 does not 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>

- e. Review the configured information and choose *Finish*.
- f. 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 \(SAP HANA Studio\) \[page 922\]](#)

## 7.2.3.2.4 Disable System Replication

To disable system replication for an SAP HANA system, you must first unregister the secondary system and then disable system replication on the primary system.

### 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.
- Both systems are online.

### 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*.

### Results

System replication is disabled.

## 7.2.3.2.5 Set Up Multitier System Replication

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.

### Prerequisites

- 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.

The secondary systems must meet the following criteria with respect to the primary system:

- They must have a different host name, or host names in the case of a distributed system.  
If the host names of the primary and the secondary systems are the same (for example, because systems that have identical host names are used) in the SAP HANA Lifecycle Manager to change the host name(s) of the secondary system.
- They must have the same number of worker hosts. This implies that if there is a standby host on the primary system it need to not be available on the secondary systems.
- They must have the same software version or higher.
- They must have the same SAP system ID (SID).

#### **i** Note

Once the SID is the same, an additional license is not required.

- They must have the same system configuration in the system properties files (\*.ini files).

#### **i** Note

Any changes made manually or by SQL on one system must be manually duplicated on the other systems. Automatic configuration parameter checks will alert you to configuration differences between the systems.

- They should run on the same endianness platform.
- The required ports must be available.  
The same instance number is used for primary, tier 2 and tier 3 secondary systems. The instance number +1 must be free on all systems, because this port range is used for system replication communications.
- You are logged on to the systems as the operating system user (user <sid>adm) or are able to enter these credentials when prompted.
- 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

Multitier system replication supports various replication mode combinations. For more information, see *Supported Replication Modes between Sites*.

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

[Supported Replication Modes between Sites \[page 840\]](#)

### 7.2.3.2.5.1 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</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
		to disk, but disk IO on tier 3 has been triggered.	
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 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 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 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</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
		<p>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>	
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</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
		<p>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 can have an impact on the primary.</p>	
ASYNC	ASYNC	<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 ASYNC 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>

### 7.2.3.3 Monitoring System Replication in the SAP HANA Cockpit

Monitoring the status of replication between the primary system and the secondary systems is important to ensure rapid takeover in the event of planned or unplanned downtime.

## 7.2.3.3.1 Monitor System Replication

To monitor system replication, you can use the [System Replication](#) app in the SAP HANA cockpit.

### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.cockpit.sysrep.roles::SysRepAdmin`. You can assign roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- The [System Replication](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it again from the tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

### Context

The [System Replication](#) app provides you with information about the status of system replication.

### Procedure

Open the [System Replication](#) app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.

The [System Replication](#) app opens displaying the [System Replication Overview](#).

### Related Information

[Assign Roles to a User \[page 717\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

#### 7.2.3.3.1.1 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"> <li>LOGREPLAY</li> <li>DELTA DATA SHIPPING</li> </ul>
Replication Mode	<p>Configured replication mode:</p> <ul style="list-style-type: none"> <li>SYNC: synchronous replication with acknowledge when buffer has been written to disk</li> <li>SYNCMEM: synchronous replication with acknowledge when buffer arrived in memory</li> <li>ASYNC: asynchronous replication</li> <li>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).</li> </ul>
Replication Status	<p>Current status of replication:</p> <ul style="list-style-type: none"> <li>UNKNOWN: secondary did not connect to primary since last restart of the primary</li> <li>INITIALIZING: initial data transfer occurs, in this state, the secondary is not usable at all</li> <li>SYNCING: secondary is syncing again (for example, after a temporary connection loss or restart of the secondary)</li> <li>ACTIVE: initialization or sync with primary is complete and secondary is continuously replicating, if crash occurs, no data loss will occur in SYNC mode</li> <li>ERROR: error occurred with the connection (details can be found in Replication Details)</li> </ul>
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"> <li>• DISABLED: full sync is not configured at all (global.ini/[system_replication]/enable_full_sync = false)</li> <li>• 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.</li> <li>• 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.</li> </ul> <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 point in time 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"> <li>• SYNC/SYNCMEM: total round trip time to send the log buffers and receive the acknowledgment.</li> <li>• 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</li> </ul>

#### 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

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

Column	Description
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

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 sync/syncmem 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 BACKLOG_SIZE since system start).
Current Replication Backlog Time (µs)	<p>Current replication backlog in microseconds, that means, the time difference between time of the last sent log buffer and the current log buffer.</p> <p>Even in replication modes sync/syncmem this column can have a value different from 0, because log buffers are still in the send queue (max number of those buffers is the number configured log buffers on primary site).</p>
Max Replication Backlog Time (µs)	Max replication backlog in microseconds (max value of BACKLOG_TIME since system startup).

## 7.2.3.3.2 Perform a Takeover

During failover, the secondary system takes over from the primary system.

### Prerequisites

- You have the privileges granted by the role `sap.hana.admin.cockpit.sysrep.roles::SysRepAdmin`. You can assign roles using the [Assign Roles](#) app of the SAP HANA cockpit. For more information, see [Assign Roles to a User](#) in the *SAP HANA Administration Guide*.
- You have the credentials of the operating system user (`<sid>adm` user) that was created when the system was installed. For more information on how to access the SAP HANA cockpit for offline administration, see [Open SAP HANA Cockpit for Offline Administration](#).
- The [System Replication](#) tile is visible on the homepage of the SAP HANA cockpit. If it's not, you can add it again from the tile catalog. For more information, see [Customize the Homepage of SAP HANA Cockpit](#).

---

## Context

The *System Replication* app provides you with information about the status of system replication.

## Procedure

1. Open the *System Replication* app by clicking the tile of the same name on the homepage of the SAP HANA cockpit.

The *System Replication* app opens displaying the *System Replication Overview*.

2. Open the *Detail* display area by clicking any row displaying information on the sites.
3. Click *Go to Secondary*.
4. Enter your user name and password.

The homepage of the SAP HANA cockpit for offline administration of the secondary system opens displaying several tiles.

5. Click the *Start, Restart, Stop* tile.

The *System Operations* view opens displaying information about the host.

6. Click *Take Over*.
7. Confirm the takeover on the secondary system by clicking *OK* in the dialog box.
8. Enter your user name and password.

The SAP HANA cockpit for offline administration indicates that the secondary system is taking over.

9. Optional: Return to the homepage of the SAP HANA cockpit for offline administration by clicking the arrow at the top of the screen.

The homepage of the SAP HANA cockpit for offline administration displays two additional tiles: SAP HANA *Documentation* and SAP HANA *Cockpit*. After the takeover took place, the secondary system is not marked as secondary anymore.

10. Optional: Open the SAP HANA Cockpit by clicking the tile of the same name on the homepage of the SAP HANA cockpit for offline administration.

After the takeover the *System Replication* tile is displayed as not configured.

## Related Information

[Open SAP HANA Cockpit for Offline Administration \[page 58\]](#)

[Assign Roles to a User \[page 717\]](#)

[Customize the Homepage of SAP HANA Cockpit \[page 27\]](#)

## 7.2.4 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 a 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.

SAP HANA clients that were configured to reach the original host need to be sent to the standby host after host auto-failover.

One approach is 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 (a multi-host list is provided in the connection string). The client connection code (ODBC/JDBC) uses a "round-robin" approach to reconnect and ensures that these clients can reach the SAP HANA database, even after failover.

Execute the following statement to get all master host candidates that should be specified in the connect string:

```
select HOST from SYS.M_LANDSCAPE_HOST_CONFIGURATION where NAMESERVER_CONFIG_ROLE like 'MASTER%' order by NAMESERVER_CONFIG_ROLE
```

To support HTTP (web) clients, which use SAP HANA Extended Application Services (SAP HANA XS), 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 the hosts on both the primary and secondary sites. For more information see, *Configuring HTTP Load Balancing for SAP HANA Extended Application Services (XS)*.

In the case that a SAP HANA instance fails the HLB, which serves as a reverse web-proxy, redirects the HTTP clients to 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.

To add an additional standby host follow the procedure as described in *Add Hosts to an SAP HANA System*.

You can monitor the status of all active and standby hosts in the Administration editor.

As of SPS 09 a Python script provides hooks that can be called in response to events during host auto-failover, see *Implementing a HA/DR Provider*.

### Related Information

[Administration Editor \[page 229\]](#)

[Configuring Clients for Failover \[page 851\]](#)

[Add Hosts Using the Command-Line Interface \[page 541\]](#)

[Configure HTTP Load Balancing for SAP HANA Extended Application Services \(XS\) \[page 853\]](#)

[Implementing a HA/DR Provider \[page 854\]](#)

## 7.2.4.1 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.

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 should be used you can execute the following SQL statement:

### Sample Code

```
select host from m_landscape_host_configuration where nameserver_config_role
like 'MASTER%' order by NAMESERVER_CONFIG_ROLE
```

One of these master candidates will be active so 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 it is replaced by the standby host. This is done because it is faster to failover to an idle standby host than 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-&gt;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>

## Client reconnect with the Secure User Store (hdbuserstore)

For the clients in a host auto-failover landscape, the use of virtual IP addresses is recommended. In the secure user store of the SAP HANA client (hdbuserstore), user logon information can be securely stored, including passwords, using the SAP NetWeaver secure store in the file system (SSFS) functionality. 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 BW) by storing a list of all (virtual) host names to which the application server can connect. All nodes that are master candidates should be added to the hdbuserstore. Please refer to SAP Note 1930853 for information on how to find out the three master candidates in a distributed system.

### Related Information

[Client Connection Recovery \[page 820\]](#)

[SAP Note SAP Note 1930853](#)

## 7.2.4.2 Configuring Application Servers for Failover

You can configure failover support for application servers by using hdbuserstore to specify a list of host names that the server can connect to.

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 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 application server will reconnect to one of the hosts specified in the host list.

You can configure the hdbclient using the following hdbuserstore command:

```
hdbuserstore SET default "<hostname_node1>:3<system_number>15; .... <hostname_node(n)>: 3<system_number>15" SAP<SID> <Password>
```

Example configuration:

```
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
```

## 7.2.4.3 Configure HTTP Load Balancing for SAP HANA Extended Application Services (XS)

In order to enable load balancing for HTTP access to SAP HANA XS, you need to set up a load balancer (for example, SAP Web Dispatcher).

### Context

To support HTTP (web) clients, which use SAP HANA Extended Application Services (SAP HANA XS), it is recommended to install an external, itself fault protected, HTTP load balancer, such as SAP Web Dispatcher, or a similar SAP-certified product from another vendor. The HTTP load balancers are configured to monitor the web-servers on all the hosts.

The SAP Web Dispatcher automatically reads the system topology from SAP HANA XS 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 1883147](#)

[Using SAP Web Dispatcher for Load Balancing with Tenant Databases \[page 182\]](#)

## 7.2.5 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.

## 7.2.5.1 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`
2. 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.

3. 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
4. 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 \[page 856\]](#)

[Install and Configure a HA/DR Provider Script \[page 861\]](#)

### 7.2.5.1.1 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
<b>startup()</b>	Beginning of name-server's start up phase	Starting nameserver	Each individual host	Nameserver aborts, start up is canceled
<b>shutdown()</b>	Just before the name-server exists	Stopping nameserver	Each individual host	Error trace is written

Name	Trigger	Caller	Landscape	Error behavior
<b>failover()</b> [Host Auto-Failover]	As soon as the nameserver 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
<b>stonith()</b> [Host Auto-Failover]	As soon as the nameserver made the decision about the new role	Master nameserver	For each failed host	Nameserver aborts, failover is canceled
<b>preTakeover()</b> [System Replication]	As soon as the hdbnsutil -sr_takeover command is issued	Master nameserver	Called only once on the master	Takeover is aborted
<b>postTakeover()</b> [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
<b>srConnection-Changed()</b>	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
<b>srServiceState-Changed()</b> [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

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')

```

```

""" 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

```

## 7.2.5.2 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.

### 7.2.5.3 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 subprocess
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 = subprocess.call(command1.split())
    self.tracer.info("command '%s' returned with rc=%s" % (command1, rc1))
    command2 = "sudo /sbin/arping -U -c 5 %s" % self.vIP
    rc2 = subprocess.call(command2.split())
    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 = subprocess.call(command.split())
    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 HADRBBase
import subprocess
class mySTONITH(HADRBBase):
    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 subprocess.call(("<script> <logon credentials> --name %s <hard
reboot> <servicing host>" % (vmName)).split())
    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
```

## 7.3 SAP HANA Database Backup and Recovery

The following sections describe how to back up and recover an SAP HANA database and how to perform periodic maintenance tasks for backups.

### **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 backup
  - Data backup
  - Storage snapshot
- Delta backup
  - Incremental backup
  - Differential backup
- Redo log backups
- Backup and recovery using third-party tools
- Integrity checks for backups
- Backup lifecycle management
- Recovery to a specific point-in-time
- Recovery to a specific data backup or storage snapshot (without using the log area or log backups)
- Database copy using backup and recovery

## Related Information

[SAP HANA Backup Types \[page 881\]](#)

[Working with Third-Party Backup Tools \[page 909\]](#)

---

[Manually Checking Whether a Recovery is Possible \[page 938\]](#)

[Copying a Database Using Backup and Recovery \[page 974\]](#)

[Recovering an SAP HANA Database \[page 949\]](#)

## 7.3.1 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 set to occur every five minutes, including during a backup. 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. With a system running on properly configured hardware, the impact on performance of savepoints is negligible.

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 \[page 264\]](#)

### 7.3.1.1 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 redo 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 redo 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.

## Related Information

[SAP HANA Backup \[page 880\]](#)

[Log Backups \[page 887\]](#)

[SAP HANA Recovery \[page 936\]](#)

[Copying a Database Using Backup and Recovery \[page 974\]](#)

## 7.3.2 Points to Note About Backup and Recovery

Before you begin preparing a backup strategy for your SAP HANA installation, you should be aware of the following information concerning backup and recovery.

You can find more information on the individual points in the subsequent sections of this documentation.

### 7.3.2.1 Points to Note: SAP HANA Backup

This section provides an overview of important points concerning SAP HANA data backups, storage snapshots, delta backups, and log backups.

- You can back up an SAP HANA database to the file system or you can use SAP-certified third-party tools. More information: *SAP HANA Backup Types* and *Working with Third-Party Backup Tools* in *Related Information*
- While full backups (data backups and storage 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.
- With a full 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.

#### **i** Note

A full data backup does not include the log area or customer-specific configuration settings.

- It is not possible to back up and recover individual database objects. Backup and recovery always apply to the whole database. More information: *Data Backups* in *Related Information*
- The content of a data backup is consistent at the start time of the 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.
- A storage snapshot captures the complete content of the SAP HANA data area at a particular point in time. More information: *Storage Snapshots* in *Related Information*
- If a new full backup (data backup or storage snapshot) is started before the previous full backup is finished, the system handles the situation as follows:
  - The first full backup continues normally
  - The second full backup does not start, and an error message is displayed

- 
- Data and logs can only be backed up when SAP HANA is online. That is, when all the configured SAP HANA services are running.
  - Data backups and log backups are created independently of each other. Differential backups and incremental backups are not created independently of each other. A differential backup cannot be created at the same time as an incremental backup.

## Related Information

[SAP HANA Backup Types \[page 881\]](#)

[Working with Third-Party Backup Tools \[page 909\]](#)

[Data Backups \[page 881\]](#)

[SAP HANA Backup \[page 880\]](#)

[Storage Snapshots \[page 886\]](#)

### 7.3.2.1.1 Points to Note: SAP HANA Multitenant Database Containers and Backup

This section provides an overview of important points concerning backup with SAP HANA multitenant database containers.

- 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 using the system database.  
This option is enabled by default.
- The backup location in the file system is specified system-wide. Backups of tenant databases are always created in sub-directories of this location.
- SAP HANA supports high isolation scenarios for third-party tools with SAP HANA multitenant database containers.  
More information: *Isolation Level High for Backups With SAP HANA Multitenant Database Containers and Third-Party Backup Tools* in *Related Information*

## Related Information

[Multitenant Database Containers \[page 15\]](#)

[Creating Backups \[page 920\]](#)

[Isolation Level High for Backups With SAP HANA Multitenant Database Containers and Third-Party Tools \[page 913\]](#)

[Converting an SAP HANA System to Support Multitenant Database Containers \[page 572\]](#)

## 7.3.2.1.2 Points to Note: Log Modes

This section provides an overview of important points concerning log modes.

After installation, SAP HANA temporarily runs in log mode `overwrite`.

In log mode `overwrite`, no log backups are created.

Running in 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

When you manually change the log mode from `overwrite` to `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.

More information: *Log Modes* and *Change Log Modes* in *Related Information*

## Related Information

[Log Modes \[page 888\]](#)

[Change Log Modes \[page 891\]](#)

## 7.3.2.1.3 Points to Note: SAP HANA Backups and Encryption

This section provides an overview of important points concerning how SAP HANA supports encryption for backups.

- SAP HANA supports encryption. However, data backups and log backups are not encrypted. If the SAP HANA database is encrypted, during backup, the data is first decrypted and then written to the data backup.
- If the SAP HANA database is encrypted, storage snapshots are also encrypted.
- You have the option to encrypt the data and log backups using third-party tools. For more information, contact your tool vendor.
- To recover an encrypted database, no special additional steps are required.

More information: *Managing Encryption of Data Volumes in the SAP HANA Database* in *Related Information*

## Related Information

[Managing Encryption of Data Volumes in the SAP HANA Database \[page 749\]](#)

## 7.3.2.1.4 Points to Note: File-Based Backups

This section provides an overview of important points concerning backing up SAP HANA to a file system.

- The configured destination for data and log backups must be valid throughout the whole system, not only for specific hosts.
- It is strongly recommended to use shared backup storage to make the backup area available to all the nodes in a database.  
Shared backup storage allows the master name server 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

[Destination for Data Backups \[page 882\]](#)

[Naming Conventions for Backups \[page 895\]](#)

[Backing Up the Backup Catalog \[page 904\]](#)

## 7.3.2.2 Points to Note: SAP HANA Recovery

This section provides an overview of important points concerning recovery of SAP HANA databases.

- 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.
- An SAP HANA database cannot be recovered to an SAP HANA database with a lower software version.  
The SAP HANA software version used for the recovery must always be the same version or higher than that of the SAP HANA database used to create the data backup or storage snapshot.  
More information: SAP Note 1948334 (*SAP HANA Database Update Paths for Maintenance Revisions*) in *Related Information*

#### **i** Note

If the backup of the source database was created using SAP HANA lower than revision 45, the backup catalog for the source database must be rebuilt.

More information: SAP Note 1812980 (*Changes to Backup Catalog with Revision 45*) in *Related Information*

- To recover the database, at least one full backup (data backup or storage snapshot) must be available before the recovery is started.  
The backup can be file-based or it can be created using a third-party tool.

#### **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, but without using log backups. A recovery to a point in time is not possible if the full backup is not in the backup catalog.

- To recover the database to a particular point in time, a full backup (data backup or a storage snapshot) must be available, together with all the relevant delta backups and log backups.
- At the beginning of a recovery, all the data and log backups to be used must be either accessible in the file system or available through a third-party tool.  
If you are recovering the database from a storage snapshot, the storage snapshot must be replicated in the data area.

#### ➔ 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.

More information: *Manually Checking Whether a Recovery is Possible in Related Information*

#### ⚠ Caution

##### Rebuilding the Backup Catalog

In exceptional situations 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.

If the backup catalog is rebuilt:

- It no longer contains storage snapshots or backups made using third-party tools.  
As a consequence, these storage snapshots or third-party tool backups can then no longer 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 AGAIN, 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 external backup catalog is affected; the internal SAP HANA backup catalog is not changed.  
The external backup catalog is used to recover SAP HANA.

More information: SAP Note 1812057 (*Recovery of the backup catalog with hdbbackupdiag*) in *Related Information*

## Related Information

[SAP Note 1948334](#)

[SAP Note 1812980](#)

[SAP HANA Recovery \[page 936\]](#)

[Manually Checking Whether a Recovery is Possible \[page 938\]](#)

[Canceling a Recovery \[page 958\]](#)

## 7.3.2.2.1 Points to Note: SAP HANA Multitenant Database Containers and Recovery

This section provides an overview of important points concerning recovery with SAP HANA multitenant database containers.

- The recovery of a tenant database is always initiated from the system database. If tenant databases need to be recovered, they are recovered individually, and not all together in one single operation.
- The system database is recovered in the same way as an SAP HANA single-container system. The 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.

### Note

If the system database is shut down for recovery, its tenant databases are automatically shut down as well. This means that, until the recovery of the system database is completed, all its tenant databases are not available.

While a tenant database is being recovered, the system database and any other tenant databases remain online.

- To recover a complete SAP HANA multitenant database container system, the system database needs to be recovered first, and then all the tenant databases are recovered individually.
- When an SAP HANA multitenant database container is recovered, the services needed are generated automatically in the tenant databases.
- A backup of an SAP HANA single-container system cannot be recovered directly to an SAP HANA multitenant database container. An SAP HANA single-container system can be migrated to an SAP HANA multitenant database system. After migrating an SAP HANA single-container system to an SAP HANA multitenant database container, a new backup history has to be started. To be able to recover a migrated database, immediately after the migration, you need to create a full data backup of both the system database and the tenant databases. More information: SAP Note 2096000 *SAP HANA multitenant database containers - Additional Information and Converting an SAP HANA System to Support Multitenant Database Containers* in *Related Information*
- A backup of an SAP HANA multitenant database container can only be recovered to an SAP HANA multitenant database container. A backup of a tenant database can be recovered to a different SAP HANA multitenant database container, except with third-party tools.
- With SAP HANA multitenant database containers, recovery using storage snapshots is currently not supported.
- Tenant database copy using backup and recovery with third-party tools is currently not supported. To copy a tenant database using backup and recovery, you can use file system-based backups.
- Adding or removing a service breaks the backup history. After you add or remove a service, you must immediately create a full backup in order to be able to recover the database.

## Related Information

[Multitenant Database Containers \[page 15\]](#)

[SAP HANA Recovery \[page 936\]](#)

[SAP Note 2096000](#)

[Converting an SAP HANA System to Support Multitenant Database Containers \[page 572\]](#)

### 7.3.2.2.2 Points to Note: Delta Backups and Recovery

This section provides an overview of important points concerning delta backups and SAP HANA recovery.

- 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 studio.
- 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.
- 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.

## Related Information

[Delta Backups \[page 883\]](#)

### 7.3.2.2.3 Points to Note: License Key and Recovery

This section provides an overview of important points concerning license key requirements for the recovery of SAP HANA databases.

The license key for an SAP HANA database is based on the system ID and the landscape ID. After a recovery, an SAP HANA license key becomes invalid if the SID or landscape ID has changed.

During recovery, a temporary license key is installed automatically if the backup used for recovery had a permanent license. 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

A license key is only needed for an SAP HANA single-container system and the system database in an SAP HANA multitenant database container. Tenant databases in an SAP HANA multitenant database container do not need a license key.

### **i** Note

If the backup that was used for recovery only had a temporary license, the database is locked immediately after recovery.

## **Related Information**

[Recovering an SAP HANA Database \[page 949\]](#)

[Prerequisites for Copying a Database Using Backup and Recovery \[page 975\]](#)

### **7.3.2.3 Points to Note: Copying a Database Using Backup and Recovery**

This section provides an overview of important points concerning copying an SAP HANA database using backup and recovery.

- It is **not possible** to copy an SAP HANA single-container system to an SAP HANA multitenant database container.
- Both the data backups and the log backups must be from either only a third-party tool or only the file system.  
For a database copy, it is not possible to mix backups from the different sources.  
(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.)
- Tenant copy using third-party tools is currently not supported. To copy a tenant database, use file system backups instead.

## **Related Information**

[Copying a Database Using Backup and Recovery \[page 974\]](#)

### **7.3.2.4 Points to Note: System Replication**

This section provides an overview of important points concerning system replication.

- Data backups and log backups can only be written on the primary system.  
The secondary system cannot write data backups or log backups.  
The secondary system only writes backups after a takeover has been completed. That is, after it has been made the new primary system.

- **⚠ Caution**

After a takeover, deactivate scheduled data backups and automatic log backups in the former primary system.

If data backups were scheduled in the original primary system, after a takeover, the data backups are scheduled to run in the new primary system with the same configuration as in the original primary system. If automatic log backups were configured, after a takeover, the log backups are created on both the new primary system and the old primary system.

- After a takeover, ensure that any backups scheduled in the new primary system are configured in accordance with your requirements.

More information: *Schedule Data Backups (SAP HANA Cockpit)* in *Related Information*

- **⚠ Caution**

After a takeover, ensure that the original primary system does not continue to write backups to the same backup destination as the new primary system.

If data backups and log backups are written to a shared network location, this location could be mounted on both the primary system and the secondary system. However, after a takeover, the original primary system still writes backups until it is stopped or until scheduled data backups or automatic log backups are disabled. As a result, the backup catalog in the shared network location may include backups from both the original primary system and the new primary system.

If backups from different systems are mixed up, a database recovery is not possible.

- After a takeover, no delta backups are allowed in the new primary system until a full data backup (data backup or storage snapshot) has been created.

- **⚠ Caution**

Disable FULL SYNC Option Before Recovery

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
```

More information: SAP Note 2165547 (FAQ: *SAP HANA Database Backup & Recovery in an SAP HANA System Replication Landscape*) in *Related Information*

- If backups are managed using a third-party tool, the `Backint` for SAP HANA API must be accessed by both the new primary system and the original primary system.

## Related Information

[Recovery with System Replication \[page 970\]](#)

[SAP Note 2165547](#)

[Schedule Data Backups \(SAP HANA Cockpit\) \[page 933\]](#)

## 7.3.2.5 Points to Note: SAP HANA on IBM Power Systems

This section provides an overview of important points concerning SAP HANA and IBM Power systems.

- It is only possible to recover SAP HANA using backups created with the same system architecture. SAP HANA backups created on an IBM Power system cannot be used to recover SAP HANA to an Intel-based system.
- If a third-party tool is certified for Intel platforms, that tool is not automatically certified for IBM Power Systems (and vice versa). Separate certification processes are required for each platform and tool version.

## 7.3.3 Authorizations for Backup and Recovery

Backup and recovery operations can only be performed by users that have the appropriate authorizations. In SAP HANA multitenant database containers, what authorization is required depends on whether administrative tasks are performed at system level or at database level.

To administer the system database of an SAP HANA multitenant database container or a single-container system, the following authorizations are required:

Task	Required Authorization
Back up SAP HANA using the Backup Console in SAP HANA studio	<ul style="list-style-type: none"><li>• BACKUP ADMIN</li><li>• CATALOG READ</li></ul>
Back up SAP HANA using SAP HANA cockpit	<p>To create backups and monitor the backup status, the privileges granted by one of the following roles is required:</p> <ul style="list-style-type: none"><li>• <code>sap.hana.backup.roles::Operator</code></li><li>• <code>sap.hana.backup.roles::Administrator</code></li></ul> <p>To schedule backups, the privileges granted by the following role is required:</p> <code>sap.hana.backup.roles::Scheduler</code>
Back up the database without a user interface	BACKUP ADMIN or BACKUP OPERATOR (recommended for batch users only)
Recover the database without a user interface	<p>This is supported for an SAP HANA single container system or the system database in an SAP HANA multitenant database container.</p> <p>The recovery is executed as the operating system user (<code>&lt;sid&gt;adm</code>). You therefore require the logon credentials of this user.</p>
Physically delete data and log backups as well as obsolete versions of the backup catalog from the backup location.	BACKUP ADMIN

Task	Required Authorization
Administration tasks executed on a tenant database through the system database	DATABASE ADMIN

### ➔ 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.

### Related Information

[Operating System User <sid>adm \[page 647\]](#)

[SAP HANA Security Guide](#)

## 7.3.4 SAP HANA Backup

There are different methods and tools to back up an SAP HANA database.

The sections that follow describe:

- The backup types supported by SAP HANA
- Redo log backups
- Backing up customer-specific configuration settings
- The backup catalog
- Backup storage in the file system and with third-party backup tools
- Creating an SAP HANA backup
- Checking the integrity of backups
- Backup audit actions for security

---

## Related Information

[Creating Backups \[page 920\]](#)

[Backing Up Customer-Specific Configuration Settings \[page 901\]](#)

[Manually Checking Whether a Recovery is Possible \[page 938\]](#)

[Naming Conventions for Backups \[page 895\]](#)

[Persistent Data Storage in the SAP HANA Database \[page 264\]](#)

### 7.3.4.1 SAP HANA Backup Types

The following sections describe the different backup types supported by SAP HANA.

#### 7.3.4.1.1 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.

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. 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 \[page 920\]](#)

[Manually Checking Whether a Recovery is Possible \[page 938\]](#)

[Backing Up Customer-Specific Configuration Settings \[page 901\]](#)

## 7.3.4.1.1.1 Destination for Data Backups

Different data backups can be written to different destinations. However, all the parts of one particular data backup are written to the same destination.

For file-based data backups, you can change the default destination. For backups made using third-party backup tools, the default backup destination cannot be changed.

### 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 from the default only for the current backup.

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.

To change the default destination for file-based data backups:

1. In SAP HANA studio, open the Backup Console for the system to be backed up.
2. Go to the *Configuration* tab.
3. In the *File-Based Data Backup Settings* area, specify the new default destination.

#### ➔ 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.

4. Save.

#### i Note

Changes to the backup destination take effect immediately.

## Destinations of Backups With Third-Party Backup Tools

Backups made using third-party tools always use the destination `/usr/sap/<SID>/SYS/global/hdb/backint`.

For third-party backup tools with SAP HANA multitenant database containers, the following directories are used:

- `/usr/sap/<SID>/SYS/global/hdb/backint/SYSTEM`
- `/usr/sap/<SID>/SYS/global/hdb/backint/DB_<tenant_database_name>`

---

It is not possible to change the backup destination for third-party tools. For this destination, the only objects created in the file system are named pipes. Named pipes occupy no space in the file system.

## Related Information

[Naming Conventions for Backups \[page 895\]](#)

[Backup Console \[page 917\]](#)

### 7.3.4.1.2 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 with 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.

#### What are Delta Backups?

Delta backups back up only data that has been changed since the last full data backup (complete data backup or storage snapshot) or the last delta backup.

##### Note

Delta backups can only be created if there is a full data backup that is delta backup-enabled.

Delta backups are supported with SAP HANA SPS10 or newer.

For 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.

### ➔ Tip

Consider partitioning column store data with a view to keeping data that is frequently changed separate from data that is not frequently changed. In this way, partitioning can help to reduce the size of delta backups.

## Related Information

[Storage Snapshots \[page 886\]](#)

### 7.3.4.1.2.1 Delta Backup Types

SAP HANA supports both differential backups and incremental backups.

Comparison of Delta Backup Types

	Differential Backup	Incremental Backup
<b>What Data is Backed Up?</b>	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).
<b>Backup Size</b>	The amount of data to be saved with each differential backup <b>increases</b> .	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.
<b>Backup and Recovery Strategy</b>	If your backup strategy is based on only full data backups and differential backups, only two backups are needed for a recovery: a full data backup and one differential backup.	If your backup strategy is based on only full data backups and incremental backups, to recover the database, SAP HANA needs the following backups: <ul style="list-style-type: none"><li>• The full data backup on which the incremental backups are based</li><li>• Each incremental backup made since the full data backup</li></ul> In some situations, a large number of incremental backups may be needed for a recovery.

### i Note

To recover SAP HANA, you can combine differential backups and incremental backups.

You can use multiple incremental backups, but only one differential backup.

### 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

## 7.3.4.1.2.2 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, you need to ensure that the backup agent executable (**hdbbackint**) is configured correctly. For information about configuring your third-party tool, consult the vendor documentation.

More information: *Configure a Third-Party Backup Tool in Related Information*

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.

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 \[page 910\]](#)

---

### 7.3.4.1.2.3 Delta Backups and the Backup Catalog

Delta backups are included in the backup catalog, but in SAP HANA studio and SAP HANA cockpit are hidden by default.

To display delta backups in SAP HANA studio:

1. In the *Systems* view, open the *Backup Console*.
2. Go to the *Backup Catalog* tab.
3. Select *Show Delta Backups*.

#### Related Information

[Display Information About Backups \(SAP HANA Cockpit\) \[page 928\]](#)  
[Backup Catalog \[page 902\]](#)

### 7.3.4.1.3 Storage Snapshots

A storage snapshot captures the content of the SAP HANA data area at a particular point in time. A storage snapshot includes all the data that is required to recover the database to a consistent state.

Storage snapshots offer an additional option to safeguard the SAP HANA data area and to recover an SAP HANA database.

Storage snapshots have the following benefits:

- Storage snapshots can be created with minimal impact on the system.  
This is because storage snapshots are created in the storage system and do not consume database resources.
- Recovery from a storage snapshot is faster than a recovery from a data backup.  
The storage snapshot only needs to be made available in the data area of the storage system. For a recovery based on a storage snapshot, you can optionally also use delta backups and log backups in the same way as with a recovery based on a data backup.

#### Storage Snapshots and Encryption

If an SAP HANA database is encrypted, storage snapshots of that database are also encrypted.

No special additional steps are required to recover an SAP HANA database from an encrypted storage snapshot.

More information: *Managing Encryption of Data Volumes in the SAP HANA Database* in *Related Information*

## Storage 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 storage snapshot. Multiple data volumes do not need to be copied synchronously; data volumes can be copied one at a time.

### Related Information

[Create a Storage Snapshot \(SAP HANA Studio\) \[page 929\]](#)

[Managing Encryption of Data Volumes in the SAP HANA Database \[page 749\]](#)

### 7.3.4.1.3.1 Storage Snapshots and Database Snapshots

A **storage 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 storage snapshot. This is particularly important if multiple storage volume groups are involved.

#### **i** Note

The internal database snapshot reflects a consistent state. While a **storage snapshot** is being created (based on the **internal database snapshot**), no further data integrity checks are performed.

(With data backups, the integrity of the data to be backed up is checked automatically while the backups are being created.)

#### **i** Note

##### **Internal Database Snapshot and System Replication**

An internal database snapshot used for a storage 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.

### 7.3.4.1.4 Log Backups

By default, SAP HANA creates redo log backups automatically at regular intervals.

During a log backup, only the actual data (the "payload") of the log segments for each service with persistence is written from the log area to service-specific log backups in the file system or to a third-party backup tool.

After a system failure, you may need log backups to recover the database to the desired state.

### Note

If the log backup area becomes temporarily unavailable, once the log backup area is available again, SAP HANA automatically continues creating log backups of all the log segments that have not so far been backed up.

### 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.

### Note

If you need to remove a service, use the procedure described in the section *Remove a Service Before or After Database Copy*, as this ensures that the log area is backed up and can be used to recover the database.

More information: *Remove a Service Before or After Database Copy* in *Related Information*

## Related Information

[Remove a Service Before or After a Database Copy \[page 988\]](#)

### 7.3.4.1.4.1 Log Modes

SAP HANA can run either in log mode `normal` or `overwrite`. This section describes the differences between these log modes.

After installation, SAP HANA temporarily runs in log mode `overwrite`.

In log mode `overwrite`, no log backups are created.

Running in 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

When you manually change the log mode from `overwrite` to `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
normal (Default)	<p>In log mode <code>normal</code>, log segments are backed up automatically if the option for automatic log backups is enabled.</p> <p>More information: <i>Enable or Disable Automatic Log Backup in Related Information</i></p> <div data-bbox="619 539 1401 689" style="background-color: #fff9c4; padding: 10px;"> <p><b>→ Tip</b></p> <p>Log mode <code>normal</code> is recommended to provide support for point-in-time recovery.</p> </div> <p>In log mode <code>normal</code>, log segments are backed up automatically in the following situations:</p> <ul style="list-style-type: none"> <li>• The log segment is full. If log segments become full, they are backed up immediately, even if the log backup interval has not been reached.</li> <li>• The log segment is closed after the configured time interval has been reached if at least one commit has been made.</li> <li>• The database is started.</li> </ul> <p>After a log segment has been backed up, SAP HANA can overwrite the space in the log area that that log segment occupied with new log entries. In this way, automatic log backups can prevent the log area from filling.</p> <p>A log segment in the log area is freed and its space can be reused only after:</p> <ul style="list-style-type: none"> <li>• The log segment is closed and a savepoint has been written.</li> <li>• The log segment has been backed up.</li> </ul> <div data-bbox="619 1279 1401 1485" style="background-color: #fff9c4; padding: 10px;"> <p><b>→ Remember</b></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> </div> <div data-bbox="619 1503 1401 1675" style="background-color: #fff9c4; padding: 10px;"> <p><b>⚠ Caution</b></p> <p>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.</p> </div>

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 <a href="#">Log Backup Settings</a> 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="630 611 1394 947" style="background-color: #fff9c4; padding: 10px;"> <p><b>⚠ Caution</b></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: <a href="#">Recover the database to a specific data backup or storage snapshot</a>.</p> </div> <div data-bbox="630 958 1394 1108" style="background-color: #fff9c4; padding: 10px;"> <p><b>i Note</b></p> <p>Even when no log backups are written, with each data backup, the <b>backup catalog</b> is still backed up to the log backup destination.</p> </div>

## Related Information

[Change a System Property \[page 217\]](#)

[Savepoints and Redo Logs \[page 869\]](#)

[Enable or Disable Automatic Log Backup \[page 894\]](#)

[Creating Backups \[page 920\]](#)

## 7.3.4.1.4.1.1 Change Log Modes

You can switch between the SAP HANA log modes: `normal` and `overwrite`

### Procedure

#### → Tip

When you manually change the log mode from `overwrite` to `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. In SAP HANA studio, go to the *Configuration* tab.
2. Open `global.ini`, then open the `persistence` section.
3. Locate the parameter `log_mode`.
4. Double-click to open the change dialog.
5. Specify the new log mode.

The new log mode can be either `normal` or `overwrite`.

To reset SAP HANA to log mode `normal`, choose *Restore Default*.

6. Save.

The change takes effect immediately.

### Related Information

[Creating Backups \[page 920\]](#)

## 7.3.4.1.4.2 Change the Destination for the Log Backups

You can change the default destination for file-based log backups.

### Context

By default, file-based log backups are written to: `$DIR_INSTANCE/backup/log`

## Procedure

### **i** Note

The default backup destination can only be changed for file-based backups.

Backups made using third-party tools always use the destination: `/usr/sap/<SID>/SYS/global/hdb/backint`

For this reason, it is not possible to change the backup destination for third-party tools.

For a destination for third-party tools, only named pipes are created in the file system. Named pipes occupy no space in the file system.

To change the default destination or the destination type:

1. In SAP HANA studio, open the Backup Console.
2. Go to the *Configuration* tab.
3. Go to **Log Backup Settings** > *Destination Type*.
4. Perform the following steps:

Task	Steps
<b>Change the destination</b>	<p>Specify the new default destination.</p> <p><b>→ Tip</b></p> <p>For improved data safety, it is recommended that you specify a path to a secure backup destination.</p> <p>The data area, log area, data backups, and log backups should never be on the same physical storage devices.</p>
<b>Change the destination type</b>	<p>Select either <i>File</i> or <i>Backint</i>.</p> <p><b>i Note</b></p> <p>The destination type <i>Backint</i> is only available if the Backint agent is installed.</p>

5. Save.

Changes to the default log backup destination take effect immediately.

### 7.3.4.1.4.3 Change the Log Backup Interval

By default, the log backup interval is 15 minutes (900s). 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`.

#### 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 of data that 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.

You can specify the log backup interval using the Backup Console in SAP HANA studio. Alternatively, you can configure the parameter `log_backup_timeout_s` in the `global.ini` configuration file.

#### Procedure

1. In SAP HANA studio, choose [Backup](#) to open the Backup Console.
2. Go to the [Configuration](#) tab.
3. Go to [Log Backup Settings](#) > [Backup Interval](#) and specify the desired log backup interval.

The default log backup interval is 15 minutes (900s). 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.

4. Save.

### Note

If log segments become full, they are backed up immediately, even if the log backup interval has not been reached.

## Results

The new log backup interval takes effect immediately.

### 7.3.4.1.4.4 Enable or Disable Automatic Log Backup

By default, automatic log backups are enabled in SAP HANA. You can manually disable or enable automatic log backups.

## Prerequisites

- To enable automatic log backups, the log mode setting must be `normal`.  
More information: *Log Modes* in *Related Information*

## Context

### Caution

During normal system operation (log mode `normal`), it is strongly recommended that you keep automatic log backup activated. If automatic log backups are disabled, the log area grows until the file system is full. If the file system becomes full, the database freezes.

## Procedure

You can disable or enable automatic log backups in SAP HANA studio.

From the Backup Console, go to the *Configuration* tab, deselect or select the option *Enable Automatic Log Backup*.

Alternatively, you can disable or enable automatic log backups in SAP HANA studio as follows:

1. From the Configuration tab:

Open  *global.ini*  *persistence section* .

- 
2. Locate the parameter `enable_auto_log_backup`.
  3. Double-click to open the change dialog.

The default setting is yes (automatic log backup is active).

4. Specify whether to enable or disable automatic log backups.

You can specify either yes or no to enable or disable automatic log backups.

5. To reset to enable automatic log backups, choose *Restore Default*.
6. 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 \[page 888\]](#)

### 7.3.4.1.5 Naming Conventions for Backups

The sections that follow describe the naming conventions and recommendations for file-based data backups, delta backups, and backups made using third-party tools.

## Related Information

[Naming Conventions for the Backup Catalog \[page 906\]](#)

#### 7.3.4.1.5.1 Naming Conventions for Data Backups

This section describes the file naming conventions 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>&lt;path&gt;</code> For example: <code>&lt;/backup/data/&gt;</code>	<p>Optional.</p> <p><b>For file-based backups:</b> if no complete path is specified, the default backup destination is prepended to the backup name.</p> <p><b>For backups made using third-party tools:</b> a named pipe is created in the file system. The named pipe is always created in the directory <code>/usr/sap/&lt;SID&gt;/SYS/global/hdb/backint</code>.</p> <p>For third-party tools with SAP HANA multitenant database containers, the following directories are used for the system database or tenant database:</p> <ul style="list-style-type: none"><li>• <code>/usr/sap/&lt;SID&gt;/SYS/global/hdb/backint/SYSTEM</code></li><li>• <code>/usr/sap/&lt;SID&gt;/SYS/global/hdb/backint/DB_&lt;tenant_database_name&gt;</code></li></ul> <p>The third-party tool reads data to be backed up from the named pipe, and writes it 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> <div data-bbox="804 488 1398 1211" style="background-color: #fff9c4; padding: 10px;"> <p><b>i Note</b></p> <p>For <b>file-based backups</b>, it is strongly recommended that you use a unique prefix for each data backup name. For example, a 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 made using <b>third-party tools</b>, 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. For example, a timestamp. A timestamp is helpful if you need to recover the database using a specific data backup without the backup catalog.</p> </div> <div data-bbox="804 1227 1398 1346" style="background-color: #fff9c4; padding: 10px;"> <p><b>i Note</b></p> <p>It is <b>not</b> possible to change the prefix.</p> </div>
<b>Suffix</b>	<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 (&lt;path&gt;&lt;prefix&gt;) 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 file 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 then 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 created with SAP HANA studio.

```
</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

## 7.3.4.1.5.2 Naming Conventions for Delta Backups

This section describes the file naming conventions for differential and incremental backups.

### Structure of File Names for Delta Backups

	Differential Backups	Incremental Backups
<b>Prefix:</b>	User-defined. A timestamp is recommended. For example: <b>2016-07-23</b>	User-defined. A timestamp is recommended. For example: <b>2016-07-23</b>
<b>String:</b>	<b>databackup_differential</b>	<b>databackup_incremental</b>
<b>Backup ID:</b>	The backup ID of the <b>full data backup</b> that the differential backup is based on	The backup ID of the <b>full data backup or the delta backup</b> that the incremental backup is based on
<b>Delta Backup ID:</b>	ID of the differential backup	ID of the incremental backup
<b>Volume ID:</b>	Volume ID as with complete data backups	Volume ID as with complete data backups
<b>Volume ID:</b>	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 ('2016-07-23')` creates a differential backup based on the previously created full data backup.

Example names of incremental backup files:

```
2016-07-23_databackup_differential_1426237023821_1426237780534_0_1
```

```
2016-07-23_databackup_differential_1426237023821_1426237780534_1_1
```

```
2016-07-23_databackup_differential_1426237023821_1426237780534_2_1
```

```
2016-07-23_databackup_differential_1426237023821_1426237780534_3_1
```

## Example

### File Names for Incremental Backups

The SQL statement `BACKUP DATA INCREMENTAL USING FILE ('2016-07-23')` creates an incremental backup based on the previously created full data backup or differential backup.

Example names of incremental backup files:

```
2016-07-23_databackup_incremental_1426237023821_1426237028496_0_1
```

```
2016-07-23_databackup_incremental_1426237023821_1426237028496_1_1
```

```
2016-07-23_databackup_incremental_1426237023821_1426237028496_2_1
```

```
2016-07-23_databackup_incremental_1426237023821_1426237028496_3_1
```

## 7.3.4.1.5.3 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
<volume ID>	The volume ID for the SAP HANA service. For example, name server, index server, script server, or XS engine.
<log partition ID>	Only one log partition is supported for each service.
<first redo log position>	The oldest entry in log backup
<last redo log position>	The most recent entry in the log backup
<backup_ID>	Uniquely identifies the log backup  <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>&lt;backup_ID&gt; is calculated automatically by SAP HANA. &lt;backup_ID&gt; is only used for file-based backups, not for backups with third-party backup tools.</p> </div>

#### Example

A log backup name could look like this:

```
<log_backup_1_0_1234567_1238567.1380740407446>
```

### 7.3.4.1.5.4 Temporary Names for File-Based Backups

When file-based backups are written, they are first written using a temporary name.

After a backup part 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**

When backups are overwritten, at least **twice the space** is needed at the backup destination, because the old backup and the new backup with the same name exist for a time in parallel.

## 7.3.4.1.6 Backing Up Customer-Specific Configuration Settings

Customer-specific configuration settings 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 using SAP HANA studio.

To display the configuration settings, go to the [Configuration](#) tab 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>\$DIR_INSTANCE/../../SYS/global/hdb/custom/config</code>
Database-specific configuration settings (For SAP HANA multitenant database containers only)	<code>\$DIR_INSTANCE/../../SYS/global/hdb/custom/config/&lt;database_name&gt;</code>
Host-specific configuration settings	<code>\$SAP_RETRIEVAL_PATH</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, these 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
sapprofile.ini	<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> <p><b>i Note</b></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> <p><b>i Note</b></p> <p><code>sapprofile.ini</code> is not displayed in the <i>Configuration</i> tab in SAP HANA studio.</p>
daemon.ini	<p>Contains information about which database services to start.</p>
nameserver.ini	<p>The <code>nameserver.ini</code> file contains global information for each installation. The landscape section contains the system-specific landscape ID and assignments of hosts to roles MASTER, WORKER, and STANDBY.</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>

### 7.3.4.1.7 Backup Catalog

The backup catalog contains information about the backup history.

The backup catalog enables SAP HANA to determine the following:

- Whether recovery is possible
- Which backups to use to recover the database
- Which backups are no longer needed

### Note

With SAP HANA multitenant database containers, the system database and each tenant database have their own backup catalog.

## Backup Catalog Contents

The backup catalog includes the following information:

- The backups created for a database  
This includes data backups, storage snapshots, delta backups (differential and incremental backups), and log backups.

### Note

Recoveries are recorded in the backup catalog, but not displayed in the monitoring views.

- The start and completion times of the backups
- Whether a backup is still running

### Note

The backup catalog does not show the progress of a backup. The progress of a backup is recorded in the `backup.log`.

- Whether a backup was successful or not
- The volumes that were backed up
- The log backups and which part of the log they contain
- Backup destinations and their sizes
- The destination type  
(file, Backint, or storage snapshot)
- A backup ID  
If you are working with a third-party backup tool, an external backup ID (EBID) is also included.

### Caution

#### Rebuilding the Backup Catalog

In exceptional situations 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.

If the backup catalog is rebuilt:

- It no longer contains storage snapshots or backups made using third-party tools.  
As a consequence, these storage snapshots or third-party tool backups can then no longer 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 AGAIN, 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 external backup catalog is affected; the internal SAP HANA backup catalog is not changed. The external backup catalog is used to recover SAP HANA.

More information: SAP Note 1812057 (*Recovery of the backup catalog with hdbbackupdiag*) in *Related Information*

## Related Information

[SAP Note 1812057](#)

[Log Files for Backup and Recovery \[page 908\]](#)

[Backing Up the Backup Catalog \[page 904\]](#)

[Monitoring Views for the Backup Catalog \[page 907\]](#)

[Housekeeping for Backup Catalog and Backup Storage \[page 990\]](#)

### 7.3.4.1.7.1 Backing Up the Backup Catalog

Each time a backup of any type is done, the backup catalog is backed up and versioned.

For file-based backups, the backup catalog is backed up to the location where the log backups are stored. This allows the backups of the backup catalog to be accessed during a recovery. 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 versioning of the backup catalog is handled by the backup tool.

## Related Information

[Housekeeping for Backup Catalog and Backup Storage \[page 990\]](#)

[Configure a Third-Party Backup Tool \[page 910\]](#)

#### 7.3.4.1.7.1.1 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 period of 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 tools.

### 7.3.4.1.7.1.1.1 Disable Writing Accumulated Backups of the Backup Catalog

By default, writing accumulated backups of the backup catalog is enabled.

#### Procedure

To disable writing accumulated backups:

1. In SAP HANA studio, go to the *Configuration* tab.
2. Open `global.ini`, then open the `backup` section.
3. Locate the parameter `enable_accumulated_catalog_backup`.
4. Double-click to open the change dialog.

The default setting is `true` (multiple log backups are accumulated to one backup of the backup catalog).

5. To back up the backup catalog after each log backup, set the parameter to `false`.

To reset to the default behavior, choose *Restore Default*.

6. Save.

### 7.3.4.1.7.1.2 Log Backups and Third-Party Tools

To improve the performance of log backups, SAP HANA uses a single backup call to the third-party agent for all the log segments of a service that are ready to be backed up at a particular time.

#### Context

If a single backup call is used for each log backup, and if a single log backup takes a long time, during that time, several other log segments may be queued for backup. During periods of high load, it can happen that log segments are closed faster than a single backup call can be executed.

In some situations, this can result in a delay in releasing log segments while they are waiting to be backed up. This can cause the log area to grow. If the log segments cannot be backed up faster than the log area is growing, the log area could even become full.

## Procedure

You can configure the number of log segments to be processed by a single backup call.

1. In SAP HANA studio, go to the *Configuration* tab.
2. Open `global.ini`, then open the `backup` section.
3. Locate the parameter `max_log_backup_group_size`.
4. Double-click to open the change dialog.

The default value is **8**.

One third-party backup process is started, and creates a maximum of 8 log backups.

### **i** Note

The process does not wait for the configured number of log backups to be queued. If fewer log backups are queued, all of the queued log backups are processed.

In the SAP HANA log area, the log segments that are part of a multiple 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. Set the appropriate value.

To reset to the default value, choose *Restore Default*.

6. Save.

### **i** Note

The backup catalog is written to and backed up **once** for multiple log backups.

In the backup catalog, separate entries are maintained for a multiple 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 if none of them is still needed for recovery.

## 7.3.4.1.7.2 Naming Conventions for the Backup Catalog

Different names are assigned to the backup catalog with file-based backups and when using a third-party tool.

The backup catalog is assigned a name in the following format:

`log_backup_0_0_0_0.<BackupID>`

With a third-party tool, the backup catalog is assigned a name in the following format:

`log_backup_0_0_0_0`

### 7.3.4.1.7.3 Monitoring Views for the Backup Catalog

You can access the backup catalog using monitoring views. Monitoring views are located 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="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></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 storage 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

## Example

### Search for a Snapshot Using M\_BACKUP\_CATALOG

To search for a 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
```

## Backup Catalog in the Backup Console

You can also access the backup catalog from the Backup Console.

More information: *Backup Catalog* in *Related Information*

## Related Information

[Backup Console \[page 917\]](#)

### 7.3.4.1.8 Log 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 or rename either file as required.

## backup.log

The `backup.log` records information about data backups, log backups, and the backup catalog. It also records information about recovery operations.

You can access the `backup.log` file in the Backup Console on the [Overview](#) tab, and in the Administration editor on the [Diagnosis Files](#) tab.

## 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`.

You can access `backint.log` in the *Administration* editor on the *Diagnosis Files* tab.

## Related Information

[View Diagnosis Files in SAP HANA Cockpit \[page 463\]](#)

### 7.3.4.2 Working with Third-Party Backup Tools

In addition to backing up an SAP HANA database 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.

## Backint for SAP HANA Interface

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.

A third-party backup tool communicates with an SAP HANA database through the `Backint for SAP HANA` interface. `Backint for SAP HANA` uses named pipes to back up the database, and performs all the actions needed to manage external storage.

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.

### **i** Note

Storage snapshots are not supported by `Backint for SAP HANA`.

---

## Prerequisites for Using Third-Party Backup Tools

- The implementation of the API of a third-party backup tools using `Backint` for SAP HANA must be certified by SAP.
- You have a support contract with the tool vendor that permits you to use the tool with SAP HANA.

More information:

- SAP Note 1730932 (Using Backup Tools with Backint for SAP HANA)
- For a current overview of certified tools, go to the *Application Development Partner 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](#)

[Application Development Partner Directory](#)

### 7.3.4.2.1 Configure a Third-Party Backup Tool

The default configuration for a third-party backup tool is defined when the tool is installed. After a backup tool has been installed, you can back up and recover an SAP HANA database without making any further changes. However, you have the option to change some of the tool configuration settings.

## Prerequisites

- The backup agent has been installed and is visible in the *Configuration* tab of the Backup Console in SAP HANA studio.  
If the backup agent is not installed, you cannot change the Backint parameter files.
- The SAP HANA database expects the backup agent executable (`hdbbackint`) to be in the following path:  
`/usr/sap/<SID>/SYS/global/hdb/opt/hdbbackint`  
If the backup 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 backup agent executable.
- 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.

## Procedure

To configure whether backups are created to the file system or using a third-party backup tool:

1. In SAP HANA studio, from the *Systems* view, open the *Backup Console*.
2. Go to the *Configuration* tab.

For third-party backup tools (`Backint` for SAP HANA), the backup destination is always `/usr/sap/<SID>/SYS/global/hdb/backupint` for both data backups and log backups.

For third-party backup tools with SAP HANA multitenant database containers, the following directories are used:

- `/usr/sap/<SID>/SYS/global/hdb/backupint/SYSTEM`
- `/usr/sap/<SID>/SYS/global/hdb/backupint/DB_<tenant_database_name>`

### **i** Note

The content of the backups is not necessarily written to these directories. The backup tool decides where the content of the backups is actually written to.

The third-party backup tool only uses the backup destination to create **named pipes** to determine where the backups are written to. The named pipes occupy no space in the file system.

The backup destinations do not need to be changed; During a recovery, SAP HANA queries information about the backup destinations from the third-party backup tool.

For third-party backup tools, you can only change the log backup interval.

3. In the *Backint Settings* section, specify the `Backint` parameter file for the backup tool.

If a `Backint` agent is installed, it is displayed.

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 documentation for the third-party backup tool.

### **i** Note

You cannot change the `Backint` agent using SAP HANA studio.

Log backup using a third-party tool is enabled or disabled in SAP HANA studio in the *Configuration* tab (see above).

You can also enable or disable log backup using a third-party tool with the `log_backup_using_backupint` parameter in the `backup` section of the `global.ini` configuration file. The default setting is **false**.

### **i** Note

If you disable `Backint`, check that the destination used for file-based backups is correct.

## Related Information

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 922\]](#)

### 7.3.4.2.1.1 Backint Timeout for Log Backups

If the Backint agent does not respond for a user-defined period of time when writing log backups, SAP HANA cancels the Backint process.

The timeout is defined by the following parameter:

```
backint_response_timeout
```

The default timeout is 600s.

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

[Log Files for Backup and Recovery \[page 908\]](#)

### 7.3.4.2.1.2 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.

For example, 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.

## Change the Number of Channels for Multistreaming

To change the number of channels for multistreaming in SAP HANA:

1. In SAP HANA studio, go to the *Configuration* tab.
2. Expand ► *global.ini* ► *backup* ▾.
3. Locate the parameter `parallel_data_backup_backint_channels`.  
This parameter controls the number of channels used for multistreaming.
4. Change the parameter value.  
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.  
To stream data backups to multiple channels, you can change the value. The maximum number of channels permitted is 32 for each service.  
Backup data is then written in parallel to the specified number of channels.  
The number of multistreaming channels applies to all data backup services larger than 128GB. Data backup services smaller than 128GB always use only one channel.

### **i** Note

Each additional channel requires an IO buffer of 512MB. 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`.

5. *Save*.

## 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.

### 7.3.4.2.1.3 Isolation Level High for Backups With SAP HANA Multitenant Database Containers and Third-Party Tools

SAP HANA supports high isolation scenarios for third-party backup tools with SAP HANA multitenant database containers.

In an SAP HANA multitenant database container with a third-party backup tool, it is necessary to ensure that tenant databases cannot access the backups of other tenant databases. If your third-party backup tool does

not support this, you can set up separate Backint parameter files for each tenant database. These Backint parameter files are managed by the operating system user (`<sid>adm`), which has read and write access. The tenant databases have read access to the Backint configuration file through the tenant-specific group. The access permissions required for system and tenant databases are described in the following section.

### **i** Note

In an SAP HANA multitenant database container with many tenant databases, many Backint parameter files may be needed to ensure high isolation.

### **➔** Tip

Check with your third-party tool vendor whether any tool-specific restrictions apply.

## Related Information

[Increase the System Isolation Level \[page 105\]](#)

[Database Isolation \[page 107\]](#)

### 7.3.4.2.1.3.1 Set the Isolation Level to High for Backups with an SAP HANA Multitenant Database Container and Third-Party Tools

By default, the isolation level is low in SAP HANA multitenant database containers. You can increase the isolation level to high to ensure that one tenant database cannot access the data backups or log backups of another tenant database.

## 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 configuration file must be owned by the operating system user `<sid>adm` and group `sapsys`.

4. Assign the access permissions 600 to the configuration 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` and `log_backup_parameter_file` to specify the Backint parameter files. You can configure these parameters in the normal way using SAP HANA studio.

## For Each Tenant 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 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.

The configuration 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 configuration 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. Open SAP HANA studio.
- b. Over the tenant database, right-click to open the context menu.
- c. Choose **► Configuration and Monitoring ► Open Administration ►**.
- d. Go to the *Configuration* tab.
- e. Locate the parameter `data_backup_parameter_file`.
- f. Double-click the parameter to display its current settings.
- g. Specify the new path to the parameter file.
- h. Save.

To change the Backint parameter file for log backups, repeat this procedure for the parameter `log_backup_parameter_file`.

Alternatively, to change the Backint parameter file setting, you can execute the following SQL statement:

```
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 multitenant container system 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
```

## 7.3.4.2.2 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.  
More information: *SAP Note 2009486 (Disable SAP HANA log backups during upgrade of third-party backup tool that supports the Backint for SAP HANA interface)* in *Related Information*
2. While you are upgrading the third-party backup software:
  - Do **not** perform a data backup of the SAP HANA system.
  - Do not do any maintenance on 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](#) 

### 7.3.4.3 Backup Console

The Backup Console is the main tool in SAP HANA studio to perform backup administration and monitoring tasks.

#### 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 <i>Overview</i> tab, choose <i>Open Backup Wizard</i> .
Prepare a storage snapshot.	From the <i>Overview</i> tab, choose <i>Manage Storage Snapshot...</i>

Task	Steps
Monitor the progress of a running backup.	<p>Go to the <a href="#">Overview</a> 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 <b>Window &gt; Preferences &gt; SAP HANA &gt; Administration &gt; Backup Console &gt; Refresh Interval in Seconds</b>.</p> <p>Alternatively, you can refresh the overview manually by choosing <a href="#">Refresh</a> in the toolbar.</p>
Review the most recent successful backup.	Go to the <a href="#">Overview</a> tab.
Cancel a running backup.	<p>Go to the <a href="#">Overview</a> tab.</p> <p>If a backup is running, the option to cancel it is displayed.</p>
Display the content of the backup log.	From the <a href="#">Overview</a> tab, choose <a href="#">Open Log File</a> .
Configure the backup settings.	<p>Go to the <a href="#">Configuration</a> tab.</p> <p>For file-based backups, you can configure the destination for the data backups and the log backups.</p> <p>For file-based data backups, you specify a default destination, which can be changed in the backup dialog. For file-based log backups, the destination specified is the actual destination to which the log backups are written.</p> <p>You can set the log backup interval and disable or enable automatic log backups.</p> <div data-bbox="804 1547 1394 1787" style="background-color: #fff9c4; padding: 10px;"> <p><b>⚠ Caution</b></p> <p>During normal system operation, it is strongly recommended that you keep automatic log backup activated. If automatic log backup is disabled, the log area grows until the file system is full. If the file system is full, the database freezes.</p> </div> <p>If you are using a third-party backup tool, you can specify a <code>Backint</code> parameter file for data backup and log backups.</p>

Task	Steps
<p>Set the maximum file size for file-based backup files.</p>	<p>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.</p> <p>If the size of a data backup file for a service exceeds the specified limit, then SAP HANA splits the file into multiple smaller files.</p> <p>Go to the <i>Configuration</i> tab and select ► <i>File-Based Data Backup Settings</i> ► <i>Limit File Size</i> ►.</p> <p>You can set the maximum file size in GB or TB.</p> <p>The maximum file size applies to the data backups of all services.</p> <div data-bbox="804 846 1396 1160" style="background-color: #fff9c4; padding: 10px;"> <p><b>i Note</b></p> <p>The actual size of backup files may be smaller than the specified maximum size.</p> <p>If existing backups are overwritten by backups with the same names, twice the space in the backup destination is needed, because the old backup and the new backup with the same name exist for a time in parallel.</p> </div> <div data-bbox="804 1173 1396 1326" style="background-color: #fff9c4; padding: 10px;"> <p><b>→ Tip</b></p> <p>The maximum file size is set by the <code>global.ini</code> parameter <code>data_backup_max_chunk_size</code>.</p> </div>

Task	Steps
<p>Monitor the backup catalog.</p>	<p>Go to the <a href="#">Backup Catalog</a> tab.</p> <p>The <a href="#">Backup Catalog</a> 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 <a href="#">Backup Details</a> 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 <a href="#">Show Delta Backups</a>.</p> <p>To display log backups, select <a href="#">Show Log Backups</a>.</p> <p>The number of entries displayed is limited to 1000. You can change this setting in the Backup Console preferences (<a href="#">Window</a> &gt; <a href="#">Preferences</a> &gt; <a href="#">SAP HANA</a> &gt; <a href="#">Administration</a> &gt; <a href="#">Backup Console</a>). Note that increasing the number of catalog entries displayed can impact the performance of the Backup Console.</p> <p>From the <a href="#">Backup Catalog</a> tab, you can remove backup entries from the backup catalog, or also physically delete the backups.</p>

## Related Information

[SAP HANA Backup Types \[page 881\]](#)

[Backup Catalog \[page 902\]](#)

[Creating Backups \[page 920\]](#)

### 7.3.4.4 Creating Backups

The sections that follow describe how to create backups.

You can create backups manually using the following tools:

- SAP HANA studio
- SAP HANA cockpit
- Native SQL
- DBA Cockpit for SAP HANA

### **i** Note

You can also schedule periodic backups using 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.

More information: *DBA Cockpit for SAP HANA* in *Related Information*

## **Related Information**

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 922\]](#)

[Create Data Backups and Delta Backups \(SAP HANA Cockpit\) \[page 926\]](#)

[DBA Cockpit for SAP HANA](#)

### **7.3.4.4.1 Estimate the Space Needed in the File System for a Data Backup**

It is important to ensure that sufficient free space is available in the file system for backups. If there is not enough space, a backup will fail. For this reason, before you back up the database, you should estimate the amount of space that will be needed in the backup destination.

## **Context**

When you back up an SAP HANA database, the estimated backup size is displayed in the backup dialog in SAP HANA studio and SAP HANA cockpit.

This information is taken from table `M_BACKUP_SIZE_ESTIMATIONS`.

More information: *M\_BACKUP\_SIZE\_ESTIMATIONS* in *Related Information*

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.

More information: *BACKUP CHECK* in *Related Information*

### **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 performed, the actual backup size may be different from the estimated size.

It is therefore recommended to keep some additional free space in reserve.

## Related Information

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 922\]](#)

[Create Data Backups and Delta Backups \(SAP HANA Cockpit\) \[page 926\]](#)

[Temporary Names for File-Based Backups \[page 900\]](#)

### 7.3.4.4.2 Create Data Backups and Delta Backups (SAP HANA Studio)

Using SAP HANA studio, you can create data backups (complete data backups and storage snapshots) and delta backups (differential backups and incremental backups) of SAP HANA single-container systems and SAP HANA multitenant database containers.

#### Prerequisites

- To back up an SAP HANA single-container system, or a system database in an SAP HANA multitenant database container, you need the authorizations BACKUP ADMIN and CATALOG READ.
- To back up a tenant database in an SAP HANA multitenant database container, you need the authorization DATABASE ADMIN.
- The database is online, and all configured services are running.  
You can 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.  
More information: [Estimate the Space Needed in the File System for a Data Backup](#) in [Related Information](#)
- For a backup using a third-party backup tool, the tool is properly configured and connected to the SAP HANA system.  
More information: [Working with Third-Party Backup Tools](#) in [Related Information](#)

#### Context

##### **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 one of the following options:

Option	Description
<i>Back Up System...</i>	Visible for an SAP HANA single-container system
<i>Backup System Database...</i>	Visible for an SAP HANA multitenant database container
<i>Backup Tenant Database...</i>	Visible for an SAP HANA multitenant database container

**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
<b>Complete Data Backup</b>	A data backup includes all the data structures that are required to recover the database.
<b>Differential Data Backup</b>	Differential backups store all the data changed since the last full data backup.
<b>Incremental Data Backup</b>	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
<b>File</b>	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.
<b>Backint</b>	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.

Option	Description
	<p><b>i Note</b></p> <p>This option is only available if a third-party backup tool is installed.</p>

- Specify the backup destination.

The default backup destination is the path specified on the *Configuration* tab of the Backup Console.

For file-based backups:	<p>Ensure that there is sufficient space at the specified backup destination.</p> <p>You can change the default backup destination.</p> <p>More information: <i>Estimate the Space Needed in the File System for a Data Backup</i> in <i>Related Information</i></p>
For third-party backup tools:	<p>For third-party backup tools, the destination is always <code>/usr/sap/&lt;SID&gt;/SYS/global/hdb/backint</code>.</p> <p>You can only change the backup prefix.</p>

- 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.

- Choose *Next*.

A summary of the backup settings is displayed.

- If all settings are correct, choose *Finish*.

The backup is started.

## 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

[Backup Console \[page 917\]](#)

[Estimate the Space Needed in the File System for a Data Backup \[page 921\]](#)

[Working with Third-Party Backup Tools \[page 909\]](#)

[Data Backups \[page 881\]](#)

[Delta Backups \[page 883\]](#)

### 7.3.4.4.2.1 Canceling a Running Data Backup or a Delta Backup (SAP HANA Studio)

You can use SAP HANA studio to cancel a running data backup or a delta backup (differential or incremental) that was started using the backup wizard.

## Prerequisites

To cancel a backup that is still running, you need the system privileges BACKUP ADMIN and CATALOG READ.

### **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 writing to a file on an NFS mount does not work.

## Procedure

1. In SAP HANA studio, open the Backup Console for the system that is being backed up.

The status of currently running backups is displayed in the [Overview](#) tab.

### **i** Note

The option to cancel a backup is only available while a backup is running.

2. Choose *Cancel Backup*.

#### **i** Note

You can also cancel a running backup from the Data Backup dialog box.

## Results

The backup is canceled and you are notified of this.

If you wish, you can start a new data backup now.

#### **→** 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.

### 7.3.4.4.3 Create Data Backups and Delta Backups (SAP HANA Cockpit)

Using SAP HANA cockpit, you can create complete data backups and delta backups (differential backups and incremental backups) of SAP HANA single-container systems. With SAP HANA multitenant database containers, you can use SAP HANA cockpit to back up the system database. Each tenant database can be backed up if you are logged on directly.

## Procedure

1. In SAP HANA cockpit, open the *Data Backup* tile.  
An overview of the information from the backup catalog is displayed.
2. Choose *Create Backup*.
3. Specify the backup type:

Option	Description
<b>Complete</b>	A data backup includes all the data structures that are required to recover the database.
<b>Incremental</b>	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).
<b>Differential</b>	Differential backups store all the data changed since the last full data backup.

The estimated size of the backup is displayed. This information is taken from table `M_BACKUP_SIZE_ESTIMATIONS`.

More information: `M_BACKUP_SIZE_ESTIMATIONS` in *Related Information*

**i Note**

Currently, SAP HANA cockpit does not support creating storage snapshots (data snapshots), backup lifecycle management, or database recovery.

- Specify the backup destination type.

Option	Description
<b>File</b>	Writes the backup data to the file system.
<b>Backint</b>	Writes the backup data through a third-party backup tool.  <div data-bbox="336 824 1393 936" style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b> This option is only available if a third-party backup tool is installed.</p> </div> <p>For information about the Backint parameters, contact your tool vendor. The Backint parameters have no effect on the behavior of SAP HANA.</p>

- 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.

- Specify the backup destination.

The default backup destination is the path specified on the *Configuration* tab of the Backup Console.

Option	Description
<b>For file-based backups:</b>	Ensure that there is sufficient space at the specified backup destination.  The default backup destination can be changed as required.  More information: <i>Estimate the Space Needed for a Data Backup</i> in <i>Related Information</i>
<b>For third-party backup tools:</b>	For third-party backup tools, the destination is always <code>/usr/sap/&lt;SID&gt;/SYS/global/hdb/backint</code> .  You can only change the backup prefix.

- To start the backup, choose *Back Up*.

The progress of the backup is displayed on the *Data Backup* tile.

- To display more details of the backup progress, click the *Data Backup* tile.

---

## Related Information

[Tile Catalog: SAP HANA Backup \[page 39\]](#)

[Create Data Backups and Delta Backups \(SAP HANA Studio\) \[page 922\]](#)

[Estimate the Space Needed in the File System for a Data Backup \[page 921\]](#)

### 7.3.4.4.3.1 Display Information About Backups (SAP HANA Cockpit)

You can customize what information from the backup catalog is displayed in SAP HANA cockpit.

#### Context

The *Data Backup* tile displays the last successful backup and the status of the most recent full data backup:

- Successful
- Running  
If a backup is running, the *Data Backup* tile displays its progress.
- Prepared  
A storage snapshot (data snapshot) has been prepared, but has not been confirmed or abandoned.
- Canceled
- Failed

#### Procedure

1. To display additional details from the backup catalog, click the *Data Backup* tile.

The following information 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

2. To display additional columns or hide columns, choose *Customize*.

In the dialog box, select or deselect the columns of interest to you.

To change the order in which the columns are displayed, use the arrow buttons.

In the same way, you can also customize the details pages of the backup catalog.

3. To filter the information displayed, choose *Filter*.

A dialog box is displayed.

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> (Storage Snapshot) are selected.</p> <p>Delta backups are included in the backup catalog, but are hidden by default.</p> <div style="background-color: #fff9c4; padding: 5px;"><p><b>i Note</b></p><p>If a delta backup fails, the filter is automatically set to include 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 only backups with a particular status:</p> <ul style="list-style-type: none"><li>○ Canceled</li><li>○ Failed</li><li>○ Prepared</li><li>○ Running</li><li>○ Successful</li></ul>
<i>Start Time</i>	<p>You can display backups from a specific time range.</p>

4. Choose *OK*.
5. To display more details about a particular backup, click its row.

### 7.3.4.4.4 Create a Storage Snapshot (SAP HANA Studio)

Using SAP HANA studio, you can create a storage snapshot for an SAP HANA single-container database with one or multiple hosts. This section describes the specific steps to create a storage snapshot using SAP HANA studio.

#### Prerequisites

The SAP HANA database is online, and all the configured services are running.

To check whether the database is online, in 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*.

#### **i Note**

Currently, storage snapshots are not supported with SAP HANA multitenant database containers.

## Context

A storage snapshot consists of all the persisted data in the data area at a particular time.

A storage snapshot is created in three steps that are performed in the SAP HANA database and at storage system level:

Create a Storage Snapshot

Step to Create a Storage Snapshot	Description
<i>Prepare</i> the database for the storage snapshot.	An <b>internal database snapshot</b> is created, and reflects a consistent database state at the point in time it is created in the file system.
Create the storage snapshot.	<p>The storage snapshot is created based on the previously created internal database snapshot.</p> <p>In the storage system, you need to manually make all the files and directories from the data area available in a separate storage location.</p> <div style="background-color: #fff9c4; padding: 5px;"><p>➔ <b>Remember</b></p><p>Storage 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.</p></div>
<i>Confirm</i> or <i>Abandon</i> the storage snapshot.	<p>If the storage snapshot was successfully made available in a new storage location, you can confirm the storage snapshot.</p> <p>A confirm may not always work, and could return an error. If the confirm fails, physically delete the storage snapshot because it cannot be used for recovery.</p> <p>To ensure its consistent state, the <b>storage snapshot</b> relies on the previously created <b>database snapshot</b>. If the database, or a database service, is restarted, the <b>database snapshot</b> is lost. If the database snapshot is lost before the storage snapshot is confirmed, the storage snapshot is still written. During confirmation, the database notifies you that the storage snapshot cannot be used.</p>

## Procedure

1. In SAP HANA studio, open the *Backup Console*.
2. Choose *Manage Storage Snapshot...*

A dialog box appears.

3. Select *Prepare*.

A dialog box appears.

4. Optionally, add a comment.

This comment helps you to identify the storage snapshot in the backup catalog.

5. Choose *OK*.

The database is prepared for the storage snapshot.

An **internal database snapshot** is created, reflecting a consistent database state at the point in time it is created.

To check that the database is now prepared for a storage snapshot, go to the *Overview* tab of the *Backup Console*. Confirmation of the time and size is displayed in the section *Status of Currently Active Data Backup*.

#### Note

As long as a database-internal snapshot exists, no new data backups or new storage snapshots can be created.

Conversely, while a data backup is running, you cannot create a storage snapshot.

At this stage, all the snapshot-relevant data is only stored in the data area. To be able to use the storage snapshot for a recovery later on, you now need to create the storage snapshot. To create a storage snapshot, this data needs to be stored in a separate location.

6. In the storage system, make all the files and directories from the data area available together in a separate storage location.

To create the storage snapshot, you can use the tool provided by your storage vendor. For more information, consult the tool documentation.

#### Note

A storage 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 storage snapshot, only the data area must be used. Do not include the log area.

#### 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 storage snapshot has been created in a separate storage location, it needs to be confirmed.

7. From the Backup Console, choose *Manage Storage Snapshot...*

A dialog box appears.

8. *Confirm* or *Abandon* the storage snapshot.

Option	Description
<b>Confirm</b>	Confirm that the storage snapshot has been successfully saved to a new storage location. You can specify an external ID to identify the storage snapshot later in the storage system.
<b>Abandon</b>	If the storage snapshot cannot be created, or if confirmation fails, choose <i>Abandon</i> . Optionally, you can add a comment to explain why the storage snapshot was not successful.

### → Tip

It is strongly recommended to confirm or abandon a storage snapshot **as soon as possible after it has been created**.

While the storage 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 storage 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 storage snapshot is confirmed, the storage snapshot is still written. During confirmation, the database notifies you that the storage snapshot cannot be used.

After you have confirmed or abandoned a storage snapshot, it is recorded in the backup catalog as either successful or unsuccessful.

The database snapshot that was used to create the storage snapshot is discarded.

It is now possible to create further storage snapshots or data backups.

## Related Information

[Storage Snapshots and Database Snapshots \[page 887\]](#)

[Backup Console \[page 917\]](#)

## 7.3.4.4.5 Schedule Data Backups (SAP HANA Cockpit)

You can schedule complete data backups or delta backups to run at specific intervals.

### Prerequisites

- To schedule backups and change backup schedules, you have the privileges granted by the role `sap.hana.backup.roles::Scheduler`.
- To schedule backups, the XS Job Scheduler must be active, and a user assigned.  
More information: *Scheduling XS Jobs in Related Information*

#### **i** Note

All times specified are in UTC.

### Procedure

1. From SAP HANA cockpit, open the [SAP HANA Backup](#) tile.  
The backup catalog is displayed.
2. Choose [Schedule Backup](#).
3. Specify the backup type.

Option	Description
<b>Complete</b>	A complete data backup includes all the data that is required to recover the database to a consistent state.
<b>Incremental</b>	Stores the data changed since the last full data backup (complete data backup or storage snapshot) or the last incremental or differential backup.
<b>Differential</b>	Stores all the data changed since the last full data backup (complete data backup or storage snapshot).

#### **i** Note

Currently, scheduling storage snapshots is not supported.

More information: *SAP HANA Backup Types in Related Information*

4. Specify the [Destination Type](#).

Option	Description
<a href="#">File</a>	If you are creating a backup to the file system, select <a href="#">File</a> . If necessary, you can specify a new destination or change the default destination.

Option	Description
	The default backup destination is the path specified on the <i>Configuration</i> tab of the <i>Backup Console</i> . More information: <i>Destination for Data Backups</i> in <i>Related Information</i>
<i>Backint</i>	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> . More information: <i>Working with Third-Party Backup Tools</i> in <i>Related Information</i>

5. Specify a *Backup Prefix*.

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.

6. Specify the schedule settings.

Schedule Settings	Description
<i>Unique Schedule Name</i>	The name of the schedule must be unique within the SAP HANA system.
<i>Start of Schedule</i>	Specify the start time of the schedule. The first backup in the series will be created after this time.
<i>Recurrence</i>	Specify whether to create backups on a weekly basis or on a monthly basis.  <b>Weekly</b> <ul style="list-style-type: none"> <li>○ By default, the week starts on a Monday. You can also start the week on a Sunday.</li> <li>○ Specify the interval between backups. By default, backups are scheduled every week.</li> <li>○ Specify on which days you want the backup to be created. You can select one or more days.</li> <li>○ Specify the time of the backups.</li> </ul> <b>Monthly</b> <ul style="list-style-type: none"> <li>○ You can specify on which day of the month, and at which monthly intervals the backups run.</li> <li>○ Specify the time of the backups.</li> </ul>

7. Save.

An overview of the backup schedule is displayed.

The next backup scheduled will run at the first possible time after the start time.

### 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, the backups will not run.

Note that when SAP HANA is running again, missed backups are **not automatically rescheduled**.

## Related Information

[Scheduling XS Jobs \[page 1110\]](#)

[SAP HANA Backup Types \[page 881\]](#)

[Destination for Data Backups \[page 882\]](#)

[Working with Third-Party Backup Tools \[page 909\]](#)

## 7.3.4.4.5.1 Backup Scheduler Options (SAP HANA Cockpit)

You can display an overview of the backup schedules, and pause, reactivate, or delete scheduled backups.

### Procedure

1. From SAP HANA Cockpit, open the *SAP HANA Backup* tile.

The backup catalog is displayed.

2. Choose *Go to Schedules*.

An overview of schedules is displayed.

Option	Description
Status	Pause or reactivate a backup schedule.
Delete	Permanently delete a schedule.
Create Schedule	Set up a new backup schedule.

Select a schedule to display its settings. From here, you can also pause, reactivate, or delete the selected schedule.

## Related Information

[Schedule Data Backups \(SAP HANA Cockpit\) \[page 933\]](#)

### 7.3.4.5 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 \[page 727\]](#)

## 7.3.5 SAP HANA Recovery

This section outlines important information concerning recovery of an SAP HANA database.

An SAP HANA database can be recovered using data backups and log backups from the following combinations of database and destination types:

- 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.
- An SAP HANA database cannot be recovered to an SAP HANA database with a lower software version. The SAP HANA software version used for the recovery must always be the same version or higher than that of the SAP HANA database used to create the data backup or storage snapshot. More information: SAP Note 1948334 (*SAP HANA Database Update Paths for Maintenance Revisions*) in *Related Information*

#### **i** Note

If the backup of the source database was created using SAP HANA lower than revision 45, the backup catalog for the source database must be rebuilt.

More information: SAP Note 1812980 (*Changes to Backup Catalog with Revision 45*) in *Related Information*

- To recover the database, at least one full backup (data backup or storage snapshot) must be available before the recovery is started. The backup can be file-based or it can be created using a third-party tool.

### **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, but without using log backups. A recovery to a point in time is not possible if the full backup is not in the backup catalog.

- To recover the database to a particular point in time, a full backup (data backup or a storage snapshot) must be available, together with all the relevant delta backups and log backups.
- At the beginning of a recovery, all the data and log backups to be used must be either accessible in the file system or available through a third-party tool.

If you are recovering the database from a storage snapshot, the storage snapshot must be replicated in the data area.

### **→** 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.

More information: *Manually Checking Whether a Recovery is Possible* in *Related Information*

Database Recovery: Database and Destination Types

SAP HANA Database Type	File	Storage Snapshot	Backint
<b>SAP HANA single-container system</b>	YES	YES	YES
<b>SAP HANA multitenant database containers</b>	YES (system database; tenant databases are recovered through the system database)	NO	YES <b>i</b> Note Currently, it is not possible to copy a tenant database using Backint, because it is not possible to change the SID or the name of the tenant database.

### **i** Note

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.

For a database copy, it is not possible to mix backups from the different sources. Both the data backups and the log backups must be from either only a third-party tool or only the file system.

## Related Information

[SAP Note 1948334](#)

[SAP Note 1812980](#)

[Points to Note: License Key and Recovery \[page 876\]](#)

[Points to Note: Delta Backups and Recovery \[page 876\]](#)

[Manually Checking Whether a Recovery is Possible \[page 938\]](#)

[Canceling a Recovery \[page 958\]](#)

[Resuming an Interrupted Recovery \[page 959\]](#)

[Copying a Database Using Backup and Recovery \[page 974\]](#)

[SAP Note 1812057](#)

[Recovery Scenarios \[page 967\]](#)

[Converting an SAP HANA System to Support Multitenant Database Containers \[page 572\]](#)

### 7.3.5.1 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 manually check backups periodically, or if you suspect that they have been changed in some way since they were created.

When SAP HANA data 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 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, and whether backups have been changed or moved since they were first written.

You can use the following tools to perform manual backup checks:

- `hdbbackupcheck`  
Checks whether individual data backups and log backups have been changed since they were created. It is recommended that you use this tool periodically.
- `hdbbackupdiag`  
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. You can use this tool before you start a recovery.

#### Caution

To maintain good recovery performance, and to allow the check to be completed quickly, `hdbbackupdiag` checks only the metadata of a backup. It does not check the integrity of the backup content on block level.

In some situations, a backup may appear to be consistent, meaning that its metadata is correct, but may actually have internal errors. In such cases, we recommend that you use `hdbbackupcheck` to check for corruption in individual data or log backups.

More information: [1869119](#)

#### **i** Note

Both `hdbbackupdiag` and `hdbbackupcheck` can be used with file system backups and third-party tools.

With third-party backup tools, `hdbbackupdiag` and `hdbbackupcheck` must run in the system with the same SID to which the backups were written.

#### **i** Note

`hdbbackupdiag` and `hdbbackupcheck` cannot be used with storage snapshots.

The tools for these manual backup checks are described in the sections that follow.

## 7.3.5.1.1 Checking Specific Backups

You can use the `hdbbackupcheck` tool to check the integrity of specific data backups and log backups for both SAP HANA single container systems or for file-based SAP HANA multitenant database containers.

### Context

To check a database in an SAP HANA multitenant database container, you need to specify the name of the backup.

`hdbbackupcheck` does not support checking tenant databases with third-party tools.

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.

### Related Information

[SAP Note 1869119 \(Checking backups using hdbbackupcheck\)](#)

## 7.3.5.1.1.1 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 for an SAP HANA installation.

### 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 from the database to an external backup tool. To check a data backup, you need to start `hdbbackupcheck` for each individual part of the data backup.

#### **i** Note

If you are working with third-party storage 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 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.

Option	Description
<code>-p &lt;directory&gt;</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 &lt;directory&gt;</code> and specify the directory.
<code>&lt;backup&gt;</code>	Name of the backup file.
<code>--backintParamFile &lt;filename&gt;</code>	Specify the configuration 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 &lt;backupid&gt;</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 &lt;ebid&gt;</code>	External backup ID  If the part of the backup is located in a third-party backup tool, you need to specify the external backup ID.
<code>--dump &lt;backupfile&gt;</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.

### 7.3.5.1.1.1 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  
DatabaseID: 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: ber130052174a  
VolumeID: 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

```
Check backup content '[DATABASE_SNAPSHOT]@node[1]'
```

### Output Code

Backup content '[DATABASE\_SNAPSHOT]@node[1]' successfully checked.

Backup '/hana/shared/BNR/HDB00/backup/data/BackupTestMaster\_databackup\_1\_1' successfully checked.

### Example

Checking a data backup that was written to the file system and at some stage was corrupted:

#### Output Code

```
hdbbackupcheck backup/data/Hallo_databackup_2_1
```

```
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 --backintParamFile /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.

## 7.3.5.1.1.2 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 `$DIR_INSTANCE/exe/SAPCAR` to the directory `<targetdir>` in the target system.
3. Unpack the archive, add the directory of `hdbbackupcheck` to the environment variable `LD_LIBRARY_PATH`, and use the program as described above:
  - a. `cd <targetdir>`
  - b. `./SAPCAR -xvf <archive>`
  - c. `export LD_LIBRARY_PATH=<targetdir>:$LD_LIBRARY_PATH`
  - d. `./hdbbackupcheck -v <backup>`

#### **i** Note

By default, the log files `backupcheck.log` and `backintcheck.log` are created in the current directory. To create the log files in a different directory, start `hdbbackupcheck` with parameter `-p <directory>`.

## 7.3.5.1.2 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, and also checks whether these backups are available and whether they can be accessed.

### Context

With `hdbbackupdiag`, you can check 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 tool:

- The backup is available in the third-party tool.

### Procedure

1. Locate the directory that contains the latest backup catalog (file-based only).

This is the directory where the last log backups were written before the recovery was started.

The default directory is `$DIR_INSTANCE/backup/log`.

#### **i** Note

`hdbbackupdiag` can read the backup catalog from a different directory.

2. Start the check program using the following command: `hdbbackupdiag [options] [-d <directory>]`

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.

Option	Description
<b>-B</b>	Display the names of backups with Backint information.
<b>-v</b>	Display all available information.  For example, the SAP HANA version that was used to create a backup.
<b>-d &lt;directory&gt;</b>	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 <b>--logDirs</b> (see below) and the third-party backup tool are also searched.  <b>i Note</b> All directories must be specified as absolute paths.
<b>-c &lt;catalog&gt;</b>	Specify the name of the backup catalog.
<b>-i &lt;BackupID&gt;</b>	Specify a backup ID.  If you do not specify a backup ID, the most recent usable data backup is used.
<b>-u &lt;"YYYY-MM-DD hh:mm:ss"&gt;</b>	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.
<b>--dataDir &lt;directory&gt;</b>	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.
<b>--logDirs &lt;directories&gt;</b>	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.
<b>--useBackintForCatalog</b>	With this option, the third-party backup tool is searched for the most recent version of the backup catalog.
<b>--databaseName &lt;database&gt;</b>	Used only with SAP HANA multitenant database containers and third-party backup tools.  This option is used with <b>--useBackintForCatalog</b> to specify a tenant database or the system database:  <b>--databaseName &lt;name_of_tenant_database&gt;</b>  <b>--databaseName SYSTEMDB</b>
<b>--backintDataParamFile &lt;paramFileName&gt;</b>	Specify a parameter file to access data backups and delta backups through a third-party backup tool.
<b>--backintLogParamFile &lt;paramFileName&gt;</b>	Specify a parameter file to access log backups through a third-party backup tool.

Option	Description
	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: <a href="#">2050606</a> (Recover database from not current backup history)
<code>--generate</code>	Generate a new backup catalog  <div style="background-color: #fff9c4; padding: 10px;"> <p><b>⚠ Caution</b></p> <p>The <code>--generate</code> option is intended for exceptional situations and <b>file-based backups only</b>. This option does not support storage 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

[Checking Specific Backups \[page 939\]](#)

### 7.3.5.1.2.1 Examples of Output From `hdbbackupdiag`

The output of `hdbbackupdiag` contains the names of all the files required to recover the 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
```

## 7.3.5.2 Recovering an SAP HANA Database

The following sections describe the options for recovering an SAP HANA database.

It may be necessary to recover an SAP HANA database in the following situations:

- Disaster recovery:
  - The data area is unusable.
  - The log area is unusable.
- Fault recovery  
If a logical error occurs, the database needs to be reset to its state at a particular point in time.
- You want to create a copy of the database.

---

## Prerequisites for Database Recovery

- The SAP HANA database software is installed on the target system.
- The number and type of services is identical in both the source and the target system. There can be any number of hosts in the target system.
- You are logged on as the operating system user `<sid>adm` in the target system.
- To recover **customer-specific configuration settings**, it is recommended that you first configure the customer-specific settings before you recover the database and the replay log backups. More information: *Backing Up Customer-Specific Configuration Files* in *Related Information*

## Related Information

[Points to Note: SAP HANA Recovery \[page 873\]](#)

[Recovering a Database From the Command Line \[page 961\]](#)

[Copying a Database Using Backup and Recovery \[page 974\]](#)

[Backing Up Customer-Specific Configuration Settings \[page 901\]](#)

## 7.3.5.2.1 Recover a Database to its Most Recent State or to a Point in Time

Using SAP HANA studio, you can recover an SAP HANA database to its most recent consistent state or to an earlier point in time.

### Prerequisites

- A full backup (complete data backup or storage snapshot) from before the specified point in time
- Delta backups made since the full backup to be used
- Log backups made since the full backup to be used (Covering changes not already contained in the delta backups)
- Log area

### Procedure

1. In SAP HANA studio, open the context menu and choose:
  - For an SAP HANA single-container system, choose **Backup and Recovery** > **Recover System...** .
  - For SAP HANA multitenant database containers, choose **Backup and Recovery**. Then choose either **Recover System Database** or **Recover Tenant Database**.

**i Note**

A **system database** in an SAP HANA multitenant database container can only be recovered to its most recent state.

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>	<p>Recovers the database to as close as possible to the current time.</p> <p>This option uses the following data:</p> <ul style="list-style-type: none"><li>○ A full backup (complete data backup or storage snapshot)</li></ul> <div data-bbox="496 1032 1396 1153"><p><b>→ Tip</b></p><p>Using the most recent available full backup makes for a faster recovery.</p></div> <ul style="list-style-type: none"><li>○ Log backups made after the selected full backup</li><li>○ Delta backups made after the selected full backup</li><li>○ The log area</li></ul> <div data-bbox="448 1261 1396 1400"><p><b>i Note</b></p><p>This is the only available option for a system database in an SAP HANA multitenant database container.</p></div>
<i>Recover the database to the following point in time</i>	<p>Specify a point in time to recover the database to.</p> <div data-bbox="448 1480 1396 1758"><p><b>i Note</b></p><p>Any changes that were made after the most recent log backup will be lost 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"><li>○ <b>Recommended:</b> The last available data backup (file-based or Backint) or storage snapshot available before the specified point in time</li><li>○ Log backups and delta backups made after the full backup and up to the desired point in time</li><li>○ Log area</li></ul>

Option	Description
	<p><b>i Note</b></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> <p><b>i Note</b></p> <p>In an SAP HANA multitenant database container, this option is available for tenant databases, but <b>not</b> for the system database.</p>
<p>▶ <i>Advanced</i> ▶  <i>Recover the database to the following log position</i> ▶</p>	<p><b>i Note</b></p> <p>This recovery type is an advanced option that can be used in exceptional cases where a previous recovery failed.</p> <p>This option uses the following data:</p> <ul style="list-style-type: none"> <li>○ The most recent data backup or storage snapshot available before the specified log position</li> <li>○ Log backups or delta backups made since the data backup or storage snapshot to be used</li> <li>○ The log area</li> </ul> <p><b>i Note</b></p> <p>Currently, this option is not available for system databases in SAP HANA multitenant database containers.</p>

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.

For an SAP HANA single-container system, storage snapshots are also displayed.

8. From the backup catalog, you can select a complete data backup or a storage snapshot. You can also resume an interrupted recovery.

Option	Description
<i>Refresh</i>	If you make available a new storage 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>	<p>To ensure that a backup exists at the specified location, choose <i>Check Availability</i>.</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> <p><b>i Note</b></p> <p>The availability of Backint backups can be checked.</p>

Option	Description
	The availability of file-based backups can only be checked if shared backup storage 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.

### Note

To recover a database from a storage snapshot, the storage snapshot must be made available in the data area of the database.

A storage snapshot can only be used to recover an SAP HANA single database container. Currently, the recovery of SAP HANA multitenant database containers using storage snapshots is not supported.

9. Choose *Next*.

10. Finalize the recovery settings.

Option	Description
<b>Check Availability of Delta and Log Backups</b>	<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> <p> <b>Caution</b></p> <p>If you choose not to perform this check before the recovery starts, the check is still performed, but later in the recovery process.</p> <p>If an error is not detected until after the recovery has been started, the recovery will be interrupted.</p> <p>After a recovery has been canceled or 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.</p> <p>If you attempt to restart the database after a recovery has been interrupted, the following message is written to the nameserver trace file:</p> <pre>Cannot start the service 'nameserver' at '&lt;host:SQL Port&gt;' responsible for the volume '&lt;volume number&gt;' because of an error during recovery.</pre> <p>In this situation, you need to recover the database using a different recovery strategy.</p> <p> <b>Note</b></p> <p><b>Shared backup storage</b></p>

Option	Description
	<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>
<b>Initialize Log Area</b>	<p>If you select this option, no log entries from the log area are replayed, and the log area is initialized. <b>The content of the log area is lost.</b> The log entries from the log backups are replayed if they are needed.</p> <p> <b>Caution</b> 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"> <li>○ The log area is unusable.</li> <li>○ You are recovering the database to a different system (database copy).</li> </ul> <p>More information: <i>Related Information</i>.</p>
<b>Use Delta Backups</b>	<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>
<b>Install New License Key</b>	<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>More information: <i>License Key and Recovery</i> in <i>Related Information</i>.</p>

11. Choose *Next*.

Option	Description
<i>Show SQL Statement</i>	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

The progress of the recovery for each SAP HANA service is displayed in the dialog box. When the recovery is complete, a message confirms this, and shows the timestamp when the recovery was completed.

The SAP HANA database is online and can be used by applications.

## Related Information

[Canceling a Recovery \[page 958\]](#)

[Data Area is Unusable \(Disaster Recovery\) \[page 967\]](#)

[Log Area is Unusable \(Disaster Recovery\) \[page 968\]](#)

[Logical Error – Point-in-Time Recovery \(Fault Recovery\) \[page 969\]](#)

[Copying a Database Using Backup and Recovery \[page 974\]](#)

[Checking the Backups Required for a Recovery \[page 945\]](#)

[Starting and Stopping Distributed SAP HANA Systems Using SAPControl \[page 1004\]](#)

### 7.3.5.2.2 Recover a Database to a Specific Full Data Backup

Using SAP HANA studio, you can recover an SAP HANA tenant database or an SAP HANA single-container system from a specific full backup.

## Prerequisites

### **i** Note

A **system database** can only be recovered to its most recent state, not to a specific full backup.

A full backup (complete data backup or storage snapshot)

### **i** Note

For a recovery from a full 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 destination type (File, Backint, or storage snapshot), the location, and the prefix of the data backup.

If you are recovering an SAP HANA single-container system from a storage snapshot, ensure that the storage snapshot is available in the database.

## 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*.
  - To recover a single-container system, choose *Recover System...*
  - To recover a tenant database, choose *Recover Tenant Database*.

The Recovery Wizard opens.

Follow the on-screen instructions.

3. If prompted, enter the **<SID>adm** user and password and choose *OK*.

### Note

**<SID>adm** is the operating system user for the system database, and is not needed for tenant databases.

4. If prompted, confirm that the database can be shut down.
5. Choose *Next*.
6. Specify the following recovery type:

Option	Description
<i>Recover the database to a specific data backup or storage snapshot (for SAP HANA single-container systems) or Recover the database to a specific data backup</i>	Recover the SAP HANA database using a data backup or storage snapshot, which you specify in the next step. <div style="background-color: #fff9c4; padding: 5px; margin-top: 10px;"> <p> <b>Caution</b></p> <p>With this option, the SAP HANA database is <b>initialized</b> with the specified data backup or storage snapshot. This data backup or storage snapshot begins a new database lifecycle. Older data backups or storage snapshots are then no longer compatible with logs written after the recovery.</p> </div> <div style="background-color: #fff9c4; padding: 5px; margin-top: 10px;"> <p> <b>Caution</b></p> <p>Log entries are not replayed, neither 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 this data backup or storage snapshot will be lost.</p> </div>

7. Choose *Next*.
8. Specify the location of the backup catalog(s).
  - *Select backup from the backup catalog*  
If the backups are not in the default location, specify the new location(s), and choose *Add*.
  - *Specify backup without catalog*  
If you use a backup that is not recorded in the backup catalog, you will need to specify the backup type (File or Backint), the location of the backup, and its prefix.

### Note

The SID of the source database is only relevant for database copy using third-party backup tools.

### Note

To recover a database from a storage snapshot, the storage snapshot must be made available in the data area of the target database. A storage snapshot can only be used to recover an SAP HANA single-container system.

Currently, the recovery of SAP HANA multitenant database containers using storage snapshots is not supported.

9. Choose *Next*.

An overview of the relevant full backups is displayed.

10. Specify the full backup to use for the recovery.

Option	Description
<i>Refresh</i>	If you make available a new storage 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>	<p>To ensure that a backup exists at the specified location, choose <i>Check Availability</i>.</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> <div data-bbox="427 1084 531 1120" data-label="Section-Header"><h3> Note</h3></div> <p>The availability of Backint backups can be checked.</p> <p>The availability of file-based backups can only be checked if shared backup storage is being used.</p>

11. Choose *Next*.

12. Finalize the recovery settings.

Option	Description
<b>Initialize Log Area</b>	<p>This option is selected by default.</p> <div data-bbox="427 1503 585 1538" data-label="Section-Header"><h3> Caution</h3></div> <p>No log entries from the log area are replayed, and the log area is initialized. <b>The content of the log area is lost</b> as with a point-in-time recovery.</p>
<b>Install New License Key</b>	<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>

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.

The SAP HANA database is online and can be used by applications.

## Related Information

[Checking the Backups Required for a Recovery \[page 945\]](#)

[Backup Catalog \[page 902\]](#)

[Starting and Stopping Distributed SAP HANA Systems Using SAPControl \[page 1004\]](#)

[Canceling a Recovery \[page 958\]](#)

### 7.3.5.2.3 Canceling a Recovery

You can cancel a recovery while it is running.

#### **i** Note

Be aware that canceling a recovery makes the database state inconsistent.

SAP HANA prevents an inconsistent database from being started. To be able to work with the database again after canceling a recovery, you need to perform the recovery again.

If a recovery is interrupted, under certain circumstances the recovery can be resumed.

To cancel a recovery while it is running, choose *Cancel Recovery* from the recovery dialog in SAP HANA studio and then confirm your decision.

If you attempt to restart the database after a recovery is canceled or interrupted, 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.
```

#### **i** Note

It is possible to cancel a recovery from the command line using `recoverSys.py`.

More information: *Options for Recovery with recoverSys.py* in *Related Information*

## Related Information

[Options for Recovery with recoverSys.py \[page 963\]](#)

[Resuming an Interrupted Recovery \[page 959\]](#)

[Recovering an SAP HANA Database \[page 949\]](#)

### 7.3.5.2.4 Resuming an Interrupted Recovery

It is possible to resume an interrupted recovery, instead of repeating the entire recovery from the beginning. It is normally only necessary to resume a recovery in exceptional circumstances.

If a recovery is canceled or interrupted, an SAP HANA database cannot start. Before work can continue in the database, the recovery must be completed. In many situations, a recovery can be repeated from the beginning reasonably quickly, and it is possible to have the database running again with only minimal delay.

However, having to repeat an interrupted recovery from the beginning may sometimes cause a significant time delay. In view of this, the option to resume an interrupted recovery can save a significant amount of time, both with a very large database or a smaller database.

#### Prerequisites

- It is possible to resume a recovery that uses a full backup with (optionally) delta backups and log backups.

#### Restriction

A recovery from **only a full data backup** cannot be resumed.

If a recovery from only a full data backup is interrupted, the recovery needs to be repeated from the beginning.

- It is possible to resume a recovery with file-based backups and third-party backup tools.
- To resume a recovery from a storage snapshot, the existing storage snapshot can be used to resume the recovery.  
The storage 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.

#### Procedure

During a recovery, SAP HANA automatically defines fallback points, which mark the point after which it is possible to resume a recovery.

The earliest valid fallback point occurs after the part of the recovery from the full data backup or from a delta backup was completed successfully. The recovery can be resumed from a point during the remaining part of the recovery, without having to repeat the part of the recovery from the full data backup or the delta backup.

Fallback points are also written for log recovery.

You can define how often fallback points are written for log recovery. With SAP HANA studio, configure the parameter `log_recovery_resume_point_interval` in the `global.ini` configuration file.

The time interval you specify translates to the maximum acceptable delay in resuming the log recovery after recovery from the data backup has been completed.

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.

### Note

If you resume a log recovery from an earlier point in time than for the first recovery, the log fallback points are not used. In this situation, only the fallback point for the recovery from the data backup can be used; the log recovery needs to be repeated from the beginning.

## Status of Resume Recovery

The fallback points are recorded in `backup.log`, which indicate whether it is possible to resume a recovery.

More information about `backup.log`: *Log Files for Backup and Recovery in Related Information*

A recovery that can be resumed is also recorded in the backup catalog. In SAP HANA studio, a partially completed recovery that can be resumed is given the backup prefix **RESUME**. To resume the recovery from that backup:

1. Start the recovery from SAP HANA studio.  
More information: *Recover a Database to its Most Recent State or to a Point in Time in Related Information*
2. In the recovery dialog, select the backup with the prefix **RESUME**.
3. Follow the steps described on-screen to complete the recovery.

## SQL Syntax to Resume a Recovery

You can recover a database using SQL statements and the tool `recoverSys.py`.

You can resume an interrupted recovery by using the SQL recovery command again, appended with `USING RESUME`.

### Example

```
RECOVER DATABASE UNTIL TIMESTAMP '2015-10-22 15:00:00' USING RESUME
```

### Note

`USING RESUME` can only be used if a fallback point already exists in the database. If no fallback point exists, an error is returned.

Fallback points are recorded in the `backup.log` file.

The following example shows how fallback points appear in `backup.log`:

```
2016-03-25T14:47:25+01:00 P123456 150dd0cfd7 INFO RECOVERY fallback point of
service: nameserver, lu123456:30201, volume: 1, snapshot id: 17

2016-03-25T14:47:25+01:00 P123456 150dd0cfd7 INFO RECOVERY fallback point of
service: xsengine, lu123456:30207, volume: 2, snapshot id: 14

2016-03-25T14:47:25+01:00 P123456 150dd0cfd7 INFO RECOVERY fallback point of
service: indexserver, lu123456:30203, volume: 3, snapshot id: 16

2016-03-25T14:47:25+01:00 P123456 150dd0cfd7 INFO RECOVERY fallback point of
service: scriptserver, lu123456:30204, volume: 4, snapshot id: 5

2016-03-25T14:47:25+01:00 P123456 150dd0cfd7 INFO RECOVERY fallback point
written into /usr/sap/HHB/SYS/global/hdb/data/mnt00001/hdb00001/
fallback_databackup_0_1
```

## Related Information

[Recover a Database to its Most Recent State or to a Point in Time \[page 950\]](#)

[Log Files for Backup and Recovery \[page 908\]](#)

[Recovering a Database From the Command Line \[page 961\]](#)

### 7.3.5.2.5 Recovering a Database From the Command Line

To recover an SAP HANA database, it is strongly recommended that you use SAP HANA studio. It is also possible to recover an SAP HANA single-container system or a system database using SQL statements and the Python script `recoverSys.py`.

## Prerequisites

### Restriction

You can run `HDBSettings.sh recoverSys.py [<parameters>]` for SAP HANA single-container systems from a master node only, and for SAP HANA multitenant database containers from the system database only.

Tenant databases cannot be recovered using the command line tool. To recover a tenant database, use SAP HANA studio or a simple SQL interface such as HDBSQL.

- The database is offline.
- You are logged with the OS user `<sid>adm`.

## Context

SQL statements to recover an SAP HANA single-container system or a system database cannot be executed using normal SQL clients, such as the `SAP HANA HDBSQL` tool, or when the database is online. For this reason, the Python script `recoverSys.py` is used to pass SQL statements to SAP HANA.

## 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

Calling `recoverSys.py` on its own will not do anything.

More information about the parameters for `recoverSys.py`: *Parameters for Recovery with recoverSys.py* in *Related Information*.

## 7.3.5.2.5.1 Recover a Database From the Command Line

Recovery from the command line is based on a full backup (complete data backup or storage snapshot) in the backup catalog. You can specify a full backup or let SAP HANA decide which backups to recover from.

## Procedure

To recover an SAP HANA single-container system or a system database:

- Set the environment using `HDBSettings.sh`.
- Execute the Python script `recoverSys.py`: `HDBSettings.sh recoverSys.py [<parameters>]`.  
`recoverSys.py` shuts down the database.

More information about the SQL statements: *SQL Statements for Recovery with Command Line Tool* in *Related Information*.

## 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.

### **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 check manually whether the recovery has completed by looking at the instance status or the logs.

## 7.3.5.2.5.2 Options for Recovery with `recoverSys.py`

The default behavior of the `recoverSys.py` tool can be overridden using the options and parameters 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
<pre>--command="&lt;SQL command&gt;"</pre>	<p>Use this option to specify a recovery command.</p> <div data-bbox="805 405 1394 689" style="background-color: #fff9c4; padding: 10px;"> <p> <b>Example</b></p> <pre>HDBSettings.sh recoverSys.py -- command="RECOVER DATABASE UNTIL TIMESTAMP '2015-10-22 15:00:00'"</pre> <p>This statement performs a recovery to the database state of '2015-10-22 15:00:00'.</p> </div> <div data-bbox="805 705 1394 1303" style="background-color: #fff9c4; padding: 10px;"> <p> <b>Example</b></p> <pre>HDBSettings.sh recoverSys.py -- command="RECOVER DATABASE UNTIL TIMESTAMP '2015-10-22 15:00:00' USING LOG PATH ('/remote/backup/CHH/log') USING BACKUP_ID 1380740407446 CHECK ACCESS USING FILE"</pre> <p>This statement performs a recovery to the database state of '2015-10-22 15:00:00' based on the data backup identified by the <b>BACKUP ID</b> '1380740407446' using the log backups locate in '/remote/backup/CHH/log'.</p> <p>The statement checks the availability of the backup files before actually performing the recovery.</p> </div> <div data-bbox="805 1319 1394 1675" style="background-color: #fff9c4; padding: 10px;"> <p> <b>Example</b></p> <p>To perform a recovery on a remote host, pass the recovery command to a remote shell command. Ensure that the quotes are passed correctly.</p> <pre>ssh &lt;sid&gt;adm@&lt;remoteHost&gt; "HDBSettings.sh recoverSys.py -- command=\"RECOVER DATABASE UNTIL TIMESTAMP '2015-10-22 15:00:00'\""</pre> </div> <div data-bbox="805 1691 1394 1809" style="background-color: #fff9c4; padding: 10px;"> <p> <b>Remember</b></p> <p>The times specified are UTC times.</p> </div>

recoverSys.py Options	Description
<pre>--wait</pre>	<p>Causes the script to wait until the recovery has completed (either successfully or unsuccessfully).</p> <p>Default: The script <b>does not wait</b> 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>More information: <i>Starting and Stopping Distributed SAP HANA Systems Using sapcontrol in Related Information</i></p> <div data-bbox="805 689 1394 902" style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></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>
<pre>--password=&lt;password&gt;</pre>	<p>If authentication is necessary, you can supply a password for <code>&lt;sid&gt;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="805 1115 1394 1261" style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></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>
<pre>--timeout=&lt;time&gt;</pre>	<p>Specify a timeout for database shutdown and start.</p> <p>Default: 120s</p>
<pre>--licenseFile=&lt;file name&gt;</pre>	<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="805 1619 1394 1888" style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>An SAP HANA license key becomes invalid if the SID 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>
<pre>--semaphoreOnly</pre>	<p>For use by SAP HANA studio only.</p>

recoverSys.py Options	Description
<code>--masterOnly</code>	For use by SAP HANA studio only.
<code>--forceMaster &lt;host&gt;</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> <p><b>i Note</b></p> <p>At most, three hosts can be used as the master host. The roles of the hosts are defined through the <code>nameserver.ini</code> file. For this reason, it is not possible to use any random host as the master host.</p>
<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> <p><b>i Note</b></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>More information: <i>Canceling a Recovery in Related Information</i></p>

## Related Information

[Starting and Stopping Distributed SAP HANA Systems Using SAPControl \[page 1004\]](#)

[Canceling a Recovery \[page 958\]](#)

## 7.3.5.3 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.

### Related Information

[SAP HANA Recovery \[page 936\]](#)

### 7.3.5.3.1 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.

Once the database has been recovered successfully from a data backup or a storage snapshot, the log entries from the log backups and the log area are replayed.

#### **i** Note

Currently, recovery of SAP HANA multitenant database containers using storage snapshots is not supported.

It is also possible to recover the database using an older data backup or storage snapshot in combination with delta backups and log backups. The log backups needed for the recovery include those created **after** the data backup or storage snapshot.

More Information: SAP Note 1821207 (Determining required recovery files) in *Related Information*.

#### **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, storage snapshot (for SAP HANA single container systems only)
- Delta backups

- Log backups
- Log area

## Steps for Recovery

Recover the database to its most recent state.

## Related Information

[Recovering an SAP HANA Database \[page 949\]](#)

[SAP Note 1821207](#) 

### 7.3.5.3.2 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, the log area cannot be used for a recovery. It is only possible to recover the data from the log backups. As a consequence, any changes that were made after the most recent log backup will be lost. 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.

#### Note

Currently, recovery of SAP HANA multitenant database containers using storage snapshots is not supported.

In the Recovery dialog, you must select the *Initialize log area* option in order to prevent the recovery of entries from the unusable log area. This option initializes the log area, and the old (unusable) content of the log area is lost.

## Used for Recovery

- Data backup  
Alternatively, storage snapshot (for SAP HANA single container 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 storage snapshot, the log entries from the log backups are replayed.
2. Select the *Initialize log area* option.

## Related Information

[Recovering an SAP HANA Database \[page 949\]](#)

[Delta Backups \[page 883\]](#)

### 7.3.5.3.3 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.

#### Note

Currently, recovery of SAP HANA multitenant database containers using storage snapshots is not supported.

#### 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, a storage snapshot (for SAP HANA single container 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.

### Note

You need to specify a point in time to recover the database to. 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 \[page 949\]](#)

### 7.3.5.3.4 Recovery with System Replication

If you are using a disaster-tolerant solution with system replication, some specific recovery scenarios apply. These scenarios are described in the sections that follow.

When you plan a recovery strategy for system replication, be aware of the following important points concerning the primary and secondary systems.

- Data backups and log backups can only be written on the primary system.  
The secondary system cannot write data backups or log backups.  
The secondary system only writes backups after a takeover has been completed. That is, after it has been made the new primary system.

-  **Caution**

After a takeover, deactivate scheduled data backups and automatic log backups in the former primary system.

If data backups were scheduled in the original primary system, after a takeover, the data backups are scheduled to run in the new primary system with the same configuration as in the original primary system. If automatic log backups were configured, after a takeover, the log backups are created on both the new primary system and the old primary system.

- After a takeover, ensure that any backups scheduled in the new primary system are configured in accordance with your requirements.

More information: *Schedule Data Backups (SAP HANA Cockpit)* in *Related Information*

-  **Caution**

After a takeover, ensure that the original primary system does not continue to write backups to the same backup destination as the new primary system.

If data backups and log backups are written to a shared network location, this location could be mounted on both the primary system and the secondary system. However, after a takeover, the original primary system still writes backups until it is stopped or until scheduled data backups or automatic log backups are

disabled. As a result, the backup catalog in the shared network location may include backups from both the original primary system and the new primary system.

If backups from different systems are mixed up, a database recovery is not possible.

- After a takeover, no delta backups are allowed in the new primary system until a full data backup (data backup or storage snapshot) has been created.

- **⚠ Caution**

Disable FULL SYNC Option Before Recovery

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
```

More information: SAP Note 2165547 (FAQ: SAP HANA Database Backup & Recovery in an SAP HANA System Replication Landscape) in *Related Information*

- If backups are managed using a third-party tool, the `Backint` for SAP HANA API must be accessed by both the new primary system and the original primary system.

## Related Information

[SAP Note 2165547](#)

[Schedule Data Backups \(SAP HANA Cockpit\) \[page 933\]](#)

### 7.3.5.3.4.1 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, storage snapshot

- Delta backups
- Log backups  
The log backups that belong to 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.  
More information: See *Configure the Secondary System* in *Related Information*
4. Start the secondary system.

## Related Information

[Configure the Primary System \[page 784\]](#)

[Configure the Secondary System \[page 786\]](#)

[Recovering an SAP HANA Database \[page 949\]](#)

### 7.3.5.3.4.2 Recovery of a New Primary System After Takeover

With system replication, after takeover, it may become necessary to recover the new primary system.

#### Note

After a takeover, the original primary system can no longer be used in the role of the primary system.

#### Caution

To ensure that no data backups and no log backups are created by the original primary system, the original primary system must be stopped.

If the original primary system continues to write log backups after the takeover, these log backups will be incompatible with the log backups written by the new primary system after the takeover.

## Used for Recovery

- Data backups  
The data backup can be created either from the original primary or the new primary system.

Alternatively, storage snapshots

### **i** Note

Recovery using storage snapshots is currently not supported for SAP HANA multitenant database containers.

- Log backups  
The log backups that belong to the data backup or the 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

### **i** Note

To recover the new primary system after takeover, you can use data and log backups that were created either on the original primary or the new primary system. The data backup or snapshot, together with the corresponding log backups, must be accessible by the new primary system.

## 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 storage snapshot) and the log backups can be accessed by the new primary system.
3. Recover the new primary system.  
More information: See *Recovering an SAP HANA Database* in *Related Information*  
When the database has been successfully recovered from the data backup or storage snapshot, the log entries from the log backups are replayed.

### **i** Note

Recovering the primary system after takeover and failback works in the same way. For a recovery, the data backup or storage snapshot and the log backups from both the original primary and the new primary systems are used.

## Related Information

[Recovering an SAP HANA Database \[page 949\]](#)

## 7.3.6 Copying a Database Using Backup and Recovery

You can create a homogeneous copy of an SAP HANA database by recovering an existing database to a different database. A homogeneous database copy is a quick way to set up a cloned database, for example, for training, testing, or development.

You can copy a database in two ways:

- To the time at which the data backup or storage snapshot was created.  
Here, you use **only** a data backup or a storage snapshot from the source database.
- To a point in time after the data backup or storage snapshot was created.  
Here, you use a data backup or a storage snapshot **and** the log backups from the source database.

An SAP HANA database can be copied using data backups and log backups from the following combinations of database and destination types:

Database Recovery: Database and Destination Types

SAP HANA Database Type	File	Storage Snapshot	Backint
<b>SAP HANA single-container system</b>	YES	YES	YES
<b>SAP HANA multitenant database containers</b>	YES (system database; tenant databases are recovered through the system database)	NO	YES  <b>i Note</b>  Currently, it is not possible to copy a tenant database using Backint, because it is not possible to change the SID or the name of the tenant database.

### **i Note**

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.

For a database copy, it is not possible to mix backups from the different sources. Both the data backups and the log backups must be from either only a third-party tool or only the file system.

### **i Note**

Backup and recovery is the recommended method for copying or moving a tenant database within the same system.

It is possible to use backup and recovery to copy or move a tenant database to a different system. In this scenario, to copy or move a tenant database to a different system securely and conveniently and with near-zero downtime, consider using system replication mechanisms.

More information: *Copying and Moving Tenant Databases Between Systems* in *Related Information*

## Related Information

[Points to Note: Copying a Database Using Backup and Recovery \[page 877\]](#)

[Copying and Moving Tenant Databases Between Systems \[page 145\]](#)

### 7.3.6.1 Prerequisites for Copying a Database Using Backup and Recovery

Before you can create a copy of an SAP HANA database, some important preparations must be made.

Ensure that the following prerequisites are met:

- A data backup or a storage snapshot of the source database is available.  
If needed, the log backups from the source database must also be available.  
More information: *SAP Note 1821207 (Determining required recovery files)* in *Related Information*
- The version of the SAP HANA target database is the same or higher than the SAP HANA source database.

#### **i** Note

If the backup of the source database was created using SAP HANA lower than revision 45, the backup catalog for the source database must be rebuilt.

More information: *SAP Note 1812980 (Changes to Backup Catalog with Revision 45)*

- 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.  
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.
- The target system can have any number of hosts, provided that the number and type of services is identical in both the source system and the target system.  
If desired, and if performance limitations are acceptable, a cloned database can be set up on a platform with less memory and CPU capacity and a different number of hosts.
- The target database must have the same number and types of services as the source database.  
If the source system has 17 worker nodes with one indexserver process running on each node, a target system with one node needs 17 indexserver processes.
- To change the number of indexserver processes in the target system (and the Backint settings, if you are using a third-party tool), a database user with the system privileges SERVICE ADMIN, INIFILE ADMIN, and BACKUP ADMIN is needed in the target system.
- For an SAP HANA single-container system, you set the number and type of services in the target database using SQL. For this reason, the target database must be online.  
For an SAP HANA multitenant database container, the number and type of services in tenant databases is set automatically during recovery.

#### **i** Note

For a database copy using storage 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 an SAP HANA single container system and the system database in an SAP HANA multitenant database container, you must have the logon credentials of the operating system user (<sid>adm).  
For an SAP HANA multitenant database container, you need the authorization DATABASE ADMIN.
- To preserve customer-specific configuration, you can make the appropriate changes manually in the configuration before you start the recovery.
- A valid license key is available for the target database.  
The license key for an SAP HANA database is based on the system ID and the landscape ID. After a recovery, an SAP HANA license key becomes invalid if the SID or the landscape ID has changed. Apply to SAP to have the license from the source database transferred to a new license key.

### **i** Note

A license key is only needed for an SAP HANA single container system and the system database in an SAP HANA multitenant database container. Tenant databases in an SAP HANA multitenant database container do not need a license key.

More information: *Points to Note: License Key and Recovery in Related Information*

- **SAP HANA on IBM Power Systems**  
An SAP HANA database can be copied using backup and recovery only from one IBM Power system to another or from one Intel-based system to another.  
More information: *Points to Note: SAP HANA on IBM Power Systems in Related Information*

## Related Information

[Points to Note: Copying a Database Using Backup and Recovery \[page 877\]](#)

[SAP Note 1821207](#)

[SAP Note 1812980](#)

[Change a System Property \[page 217\]](#)

[Add or Remove Hosts from an SAP HANA System \[page 1003\]](#)

[Add Services Before a Database Copy \[page 987\]](#)

[Remove a Service Before or After a Database Copy \[page 988\]](#)

[Points to Note: License Key and Recovery \[page 876\]](#)

[Points to Note: SAP HANA on IBM Power Systems \[page 879\]](#)

## 7.3.6.2 Copy a Database to a Point in Time Using File-Based Backups or Storage Snapshots

You can create a homogeneous copy of a database using both a data backup or a storage snapshot and log backups. Using log backups allows you to recover the database to a point in time after the data backup or storage snapshot was created.

### Prerequisites

#### Note

In SAP HANA multitenant database containers, the system database and the tenant databases can only be copied using file-based backups.

Before you can create a copy of a database using file-based backups or storage snapshots, the following prerequisites must be met:

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 in the target database.

Make the data backups and the log backup files from the source database available in the appropriate directory in the target database. If you are using storage snapshots from the source database, make them available in the data area in the target database.

#### Note

If you are using a storage snapshot, you first need to stop the SAP HANA database, then make available the storage snapshot in the target database.

#### Note

If log mode OVERWRITE is active in the target database, a database copy can still be performed using log backups that were created in log mode NORMAL in the source database.

#### Caution

The content of the log area of the source database cannot be used for recovery.

If the option *Recover the database to its most recent state* or *Recover the database to the following point in time* was selected, the additional option *Initialize Log Area* is mandatory.

More information: *Related Information*

## Procedure

1. Start the recovery of the target database.

In the *Systems* view, right-click the target database and choose ► *Backup and Recovery* ► *Recover System...* .

2. Specify the recovery type.

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. Specify the location of log backups on the target database 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 storage snapshot.
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

[Steps After Copying a Database \[page 984\]](#)

## 7.3.6.3 Copy a Database Using File-Based Data Backup or Storage Snapshot Only

You can create a homogeneous copy of a database using only a data backup or a storage snapshot. The content of the copied database is exactly the same as at the point in time at which the data backup or storage snapshot was created.

### Prerequisites

#### **i** Note

In SAP HANA multitenant database containers, the system database and the tenant databases can only be copied using file-based backups.

Before you can create a copy of a database using file-based backups or storage snapshots, the following prerequisites must be met:

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 in the target database.

Make the data backups from the source database available in the appropriate directory in the target database. If you are using storage snapshots from the source database, make them available in the data area in the target database.

#### **i** Note

If you are using a storage snapshot, you first need to stop the SAP HANA database, then make available the storage snapshot in the target database.

More information: *Related Information*

### Procedure

1. Start the recovery of the target database.

In the *Systems* view, right-click the target database and choose ► *Backup and Recovery* ► *Recover System...* .

2. Specify the recovery type *Recover the database to a specific data backup or storage snapshot*.
3. Choose *Next*.
4. Specify the backup location.

You can select *Select backup from the backup catalog* and specify the location of the backup catalog, or select *Specify backup without 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 storage 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

[Steps After Copying a Database \[page 984\]](#)

### 7.3.6.4 Copying a Database Using Third-Party Tools

Using third-party tools, you can create a homogeneous copy of an SAP HANA database.

#### **i** Note

Using third-party tools, it is only possible to copy SAP HANA single container systems. It is not possible to make system copies of SAP HANA multitenant database containers.

#### **i** 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 third-party backup agent (Backint) 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.

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.

### 7.3.6.4.1 Copy a Database to a Point in Time Using Third-Party Tools

Using third-party tools with a data backup and log backups, you can create a homogeneous copy of an SAP HANA single-container system to a specific point in time.

#### Procedure

1. Start the recovery of the target database.

In the *Systems* view, right-click the target database and choose ► *Backup and Recovery* ► *Recover System...* .

2. Choose *OK* to shut down the SAP HANA database.
3. Specify the recovery type.

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*.  
By default, the *Source System* is set to the SID of the target database. Change this to the SID of the source database that you want to copy.

#### **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.

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.

### **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 \[page 984\]](#)

### 7.3.6.4.2 Copy a Database Using a Data Backup Only and Third-Party Tools

Using third-party tools with a data backup only, you can create a homogeneous copy of an SAP HANA single-container system.

## Procedure

1. Start the recovery of the target database.

In the *Systems* view, right-click the target database and choose ► *Backup and Recovery* ► *Recover System...* .

2. Choose *OK* to shut down the SAP HANA database.
3. Specify *Recover the database to a specific data backup or storage snapshot*.
4. Choose *Next*.
5. Specify the backup location.

You can select *Select backup from the backup catalog* and *Search for the catalog in Backint only*, or select *Specify backup without catalog*.

6. Specify the SID of the *Source System*.

By default, the *Source System* is set to the SID of the target database.

Change this to the SID of the source database that you want to copy.

### **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.

7. Choose *Next*.
8. Select a backup catalog or the destination type *Backint*.  
If required, specify the backup prefix.
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 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 \[page 984\]](#)

### 7.3.6.5 Steps After Copying a Database

When you have completed a database copy, consider performing the following important steps:

- **File-system:** It is not necessary to immediately create a new backup of the target database. If you need to recover the target database before you have created a full backup of the target database, you will still be able to use backups from the source database.  
Nevertheless, it is recommended to back up the target database after recovery is completed.  
Also, it is recommended to keep the old backups available, at least until a new data backup has been created.

- **Third-party tools:** When a database copy is created, a new backup catalog is created in the target database. This backup catalog allows the target database to be recovered again, if necessary, using backups from the source database and new backups from the target database.  
To recover the target database again, you need to perform a point-in-time recovery. 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.
- 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.
- After a recovery to create a database copy, the system may include different volumes, or volumes may be assigned to a different host. 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. 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.
- If the number or type of services has changed, when the database copy is completed, you should return to the normal service configuration, with only one single service of a type on a host.

### Caution

#### Scheduled Backups after a Database Copy

If backups were scheduled in the source database, after a database copy, the schedules are also active in the target database. The backups that were scheduled in the source database are now scheduled to run in the target database with the same configuration as in the source database.

After a database copy, ensure that any backups scheduled in the target database are configured in accordance with your requirements.

More information: *Schedule Data Backups (SAP HANA Cockpit)* in *Related Information*

## Related Information

[Remove a Service Before or After a Database Copy \[page 988\]](#)

[Backup Catalog \[page 902\]](#)

[Schedule Data Backups \(SAP HANA Cockpit\) \[page 933\]](#)

## 7.3.6.6 Database Copy: Scenarios

In the following example scenarios for database copy, the source database and the target database are both SAP HANA single-container systems.

### Example

#### Target Database has Fewer Hosts Than the Source Database

You are performing a database copy where the target database has fewer hosts than the source database.

The same prerequisites apply as for an SAP HANA database recovery.

More information: *Related Information*

In the following scenario, the source database has two hosts, each with one index server. The target database has only one host. You are working with file-based backups.

1. Create a target database with one host.

#### Note

For the recovery, the type and number of services in the source database backup must be the same as the type and number of services in the target database.

2. Configure an additional index server service on the target host.

More information: *Related Information*

3. Copy the database.

Follow the steps described in *Copy a Database Using File-Based Data Backup or Storage Snapshot Only*. All content of the data backup is recovered to the target host. As a result, 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.

4. Remove the superfluous index server.

More information: *Related Information*

The source database with two hosts has been copied to the target database with one host.

### Example

#### Target Database has More Hosts Than the Source Database

You are performing a database copy where the target database has more hosts than the source database.

The same prerequisites apply as for a recovery.

More information: *Related Information*

In the following scenario, source database has two hosts, each with one index server. The target database has three hosts. You are using a third-party backup tool.

1. Create a target database with two hosts.

The third host in this constellation remains unused at this stage.

2. Copy the database.

Follow the steps described in *Copy a Database Using Data Backup Only and Third-Party Backup Tools in Related Information*.

All content of the backup is recovered to the two existing hosts.

3. Add the remaining host to the target database.  
More information: *Related Information*
4. Distribute all the data in the target database from the two hosts you had to start with to the three hosts you have now.  
A database with three hosts has been created, containing the data from the previous two-host database.

## Related Information

[SAP HANA Recovery \[page 936\]](#)

[Copy a Database Using File-Based Data Backup or Storage Snapshot Only \[page 979\]](#)

[Add Hosts Using the Command-Line Interface \[page 541\]](#)

[Remove Hosts Using the Command-Line Interface \[page 549\]](#)

[Add Services Before a Database Copy \[page 987\]](#)

[Remove a Service Before or After a Database Copy \[page 988\]](#)

[Redistribution of Tables in a Distributed SAP HANA System \[page 1007\]](#)

[Copy a Database Using a Data Backup Only and Third-Party Tools \[page 983\]](#)

[Copying a Database Using Third-Party Tools \[page 980\]](#)

## 7.3.6.7 Add Services Before a Database Copy

If the target database has fewer hosts than the source database, before you start a database copy, it is necessary to add missing services to make the number and type of services the same in both the source and the target, or to maintain an equal distribution of services in scale-out instances.

### Prerequisites

- You have the system privilege SERVICE ADMIN.
- The database is online to enable the use of SQL.

### Context

It is only necessary to add services to SAP HANA single-container systems.

With SAP HANA multitenant database container, the services are generated automatically. You do not need to take any special steps to change the number of services for an SAP HANA multitenant database container. However, you can create services on specific hosts to ensure optimal distribution of services.

For SAP HANA single-container systems, you can set the type and number of services on the target database using the SQL statement `ALTER SYSTEM ALTER CONFIGURATION`.

### Example

In an SAP HANA single-container systems, to add additional index server services, use the following SQL statement:

```
ALTER SYSTEM ALTER CONFIGURATION('daemon.ini','host','<hostname>')
SET('indexserver.c','instanceids')='<n>[,<n>+2]' WITH RECONFIGURE
```

Here, the instance ID (n) is incremented by 2 for every additional service.

### Example

In a single-host database, to add three additional index server services, use the following SQL statement:

```
ALTER SYSTEM ALTER CONFIGURATION('daemon.ini','host','host1')
SET('indexserver.c','instanceids')='40,42,44' WITH RECONFIGURE
```

Here, the instance ID (n) starts with a value of 40 and is incremented by 2 for every additional service.

### Example

In a database with multiple hosts, to add one index server service to the host called host1, use the following SQL statement:

```
ALTER SYSTEM ALTER CONFIGURATION('daemon.ini','host','host1')
SET('indexserver.c','instanceids')='40' WITH RECONFIGURE
```

The target database now has the desired number and type of services.

## Related Information

[Multitenant Database Containers \[page 15\]](#)

## 7.3.6.8 Remove a Service Before or After a Database Copy

If more than one service of a particular type is running on a host, you can remove the superfluous service(s) to streamline the use of database resources.

### Prerequisites

- You have the system privileges *SERVICE ADMIN*, *CATALOG READ* and *RESOURCE ADMIN*.
- The database is *online*.
- The service to be removed is *running*.

- Automatic log backup must be enabled.
- You can only remove services that have their own persistence.  
When the service is removed, if data is still stored in the service's persistence, its data is re-distributed to other existing services.

### **i** Note

At least one **index server** must be running on each host. If no index server running on a host, that host is orphaned. For this reason, you cannot remove the primary index server on a host.

You cannot remove the name server or the master index server in an SAP HANA system.

More information: *Related Information*

## Context

You can remove a service in the following situations:

- **Before** a database copy  
It is only necessary to remove services from SAP HANA single-container systems. There is no need to remove services from an SAP HANA multitenant database container, as the services are regenerated automatically.  
For example, if the source database has too many services.
- **After** a database copy  
You may need to remove services from either an SAP HANA single-container systems or an SAP HANA multitenant database container.  
For example, after a database copy to a target database with fewer hosts and more services than the source database.

### **i** Note

Removing a service from an **SAP HANA single-container system** does not break its backup history.

However, to ensure that the backup history remains intact, you must follow the steps described below and you must not change the SAP HANA configuration.

More information about SAP HANA multitenant database containers: *Remove a Service from a Tenant Database* in *Related Information*

## Procedure

1. In SAP HANA studio, go to the *Systems* view, right-click the database and choose **Configuration and Monitoring** **Open Administration**.
2. Open the *Landscape* tab.
3. Right-click the service to be removed and choose *Remove Service*.

## Results

The service is stopped, and the service entries are removed from the configuration.

### **i** Note

To remove a **host**, use SAP HANA database lifecycle manager (HDBLCM).

More information: *Related Information*

## Related Information

[Redistribution of Tables in a Distributed SAP HANA System \[page 1007\]](#)

[Database Copy: Scenarios \[page 986\]](#)

[Add Services Before a Database Copy \[page 987\]](#)

[Operations on Services \[page 295\]](#)

[Remove a Service from a Tenant Database \[page 138\]](#)

## 7.3.7 Housekeeping for Backup Catalog and Backup Storage

To keep your backup storage space at an optimum level, you should regularly delete backups that are no longer needed for a recovery.

To free up backup storage, you can physically delete data backups and log backups, and delete their associated entries in the backup catalog. To reduce the size of the backup catalog, you can 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.

### **➔** Tip

It is important to regularly truncate the backup catalog because, as it increases in size, it takes longer to record each new backup.

### **i** Note

If a backup is physically available, but not recorded in the backup catalog, that backup can still be used to recover the database.

## Related Information

[Backup Catalog \[page 902\]](#)

---

## 7.3.7.1 Prerequisites for Deleting Old Data and Log Backups

Before you delete old backups, some important prerequisites must be met.

- You have system privilege **BACKUP ADMIN**.
- If you are deleting multiple backups, you have decided from which time onwards you want to retain the 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 Storage Snapshots

If storage snapshots are deleted from the backup catalog, they are **not deleted physically**.

Storage 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 kept in a secure location. Ensure that these backups cannot be accessed directly by SAP HANA and cannot be deleted. An archived backup can still be used to recover SAP HANA, even if it is not recorded in the backup catalog.

## 7.3.7.2 Delete Old Data and Log Backups

You can delete one or more data backups and log backups from the backup catalog only, or from the backup catalog and also physically from the backup location.

### Procedure

#### Caution

Deleting data backups is a critical operation. Use the auditing feature in the SAP HANA database to make sure that a record is kept of all backup deletions. The action to be audited is BACKUP CATALOG DELETE.

#### Note

The steps described here for truncating the backup catalog and for deleting data backups and log backup ensure that SAP HANA can still be recovered. For this reason, it is not possible to delete log backups individually.

1. In SAP HANA studio, choose **Backup and Recovery** > **Open Backup Console**.
2. In the Backup Console, go to the **Backup Catalog** tab.
3. Select a data backup and right-click to open the context menu.

Option	Description
<i>Delete Data Backup...</i>	Delete the selected data backup.
<i>Delete Older Backups...</i>	Retain the selected data backup and delete all older data backups and log backups.

A dialog box appears.

4. Specify what you want to delete:

Option	Description
<b>Catalog</b>	Delete the selected data backup(s) from the backup catalog only.
<b>Catalog and Backup Location</b>	Delete the selected data backup(s) from the backup catalog and physically from the backup location. Specify the location of the physical data backup(s) to be deleted. 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.

5. 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.

6. 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 corresponding entries and refreshes the backup catalog in the Backup Console.

If you chose to delete from the backup catalog **and the backup location**, the system first deletes the entries in the backup catalog, and then refreshes the backup catalog in the Backup Console. 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).

Physical deletion takes place asynchronously in the background. Depending on the size of the backup, 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.

### **i** Note

If the oldest data backup is deleted, the log backups for the period up to the next oldest data backup are **not** automatically deleted.

## Related Information

[Create an Audit Policy \[page 727\]](#)

[Manually Checking Whether a Recovery is Possible \[page 938\]](#)

## 7.3.8 Planning Your Backup and Recovery Strategy

When you plan a backup strategy, consider using a combination of data backups, automatic log backups, and storage snapshots to minimize the risk of data loss and to ensure that, if a problem occurs, 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

#### **i** Note

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:

- Storage snapshot: daily
- Complete data backup using File or Backint: once a week
- Automatic log backups

## Comparison of Complete Data Backups and Storage Snapshots

Below is an overview of the comparative features of complete data backups and storage snapshots. You can use this overview to weigh up the benefits of using each backup type in your backup strategy.

### Advantages and Disadvantages of Data Backups and Storage Snapshots

	Complete Data Backup to File	Complete Data Backup Using Backint	Storage Snapshot
<b>Advantages</b>	<ul style="list-style-type: none"> <li>Integrity checks at block level</li> </ul>	<ul style="list-style-type: none"> <li>Integrity checks at block level</li> <li>Integrated into existing data center infrastructure</li> <li>Backup tool offers additional features. For example, encryption or de-duplication.</li> <li>Backups are immediately available for recovery.</li> </ul>	<ul style="list-style-type: none"> <li>Fast</li> <li>Generates negligible network load</li> <li>Can be encrypted (if data volume encryption is active.)</li> </ul>
<b>Disadvantages</b>	<ul style="list-style-type: none"> <li>Requires additional storage</li> <li>Generates additional network load</li> <li>File system needs to be monitored (fill level)</li> <li>More time is needed to make backups available for recovery</li> </ul>	<ul style="list-style-type: none"> <li>Generates additional network load</li> </ul>	<ul style="list-style-type: none"> <li>No integrity checks at block level</li> </ul>
<b>Backup Size</b>	<ul style="list-style-type: none"> <li>Payload only</li> </ul>	<ul style="list-style-type: none"> <li>Payload only</li> </ul>	<ul style="list-style-type: none"> <li>Size of the data area (but is usually compressed or de-duplicated by storage tool)</li> </ul>
<b>Backup Duration</b>	<ul style="list-style-type: none"> <li>IO-bound (reading from data volume, writing to target)</li> <li>Network-bound (writing to target file system)</li> </ul>	<ul style="list-style-type: none"> <li>IO-bound (reading from data volume)</li> <li>Network-bound (writing to backup server)</li> </ul>	<ul style="list-style-type: none"> <li>Negligible (depending on the storage tool implemented)</li> </ul>

## Related Information

[Points to Note About Backup and Recovery \[page 870\]](#)

[SAP HANA Backup \[page 880\]](#)

[SAP HANA Recovery \[page 936\]](#)

## 7.3.9 Reference: Backup Alerts

This section provides an overview of alerts that warn you of errors related to data and log backups.

Alert:	Check availability of volumes for backup
Alert ID:	34
Description:	This check warns you if a backup cannot be performed because a volume or a service is unavailable.  This alert can be triggered in combination with the alerts NOT_ASSIGNED_VOLUMES and CHECK_INACTIVE_SERVICES.
Alert Text:	<no> (<service>) is not available. A backup cannot be performed.
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, perform 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:	Checks whether the last data backup was successful, and warns you if the last data backup failed.  If a scheduled backup fails, this check can help you prevent a situation from arising where no current backups are available.
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 perform a new data backup as soon as possible.
Default Interval:	24 hours

<b>Alert:</b>	<b>Check the age of the last data backup</b>
Alert ID:	37
Description:	Checks the age of the last successful data backup, and provides different levels of notification if the last successful data backup is too old.
Alert Text:	The last data backup is <days> days old.
User Action:	To reduce your downtime in a recovery situation, perform a data backup as soon as possible.
Default Interval:	24 hours

<b>Alert:</b>	<b>Check last log backups.</b>
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 performed automatically, this is the only way to notify users. This check should therefore be performed frequently and be accorded high priority.</p> <div style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;"> <p><b>i Note</b></p> <p>Log backups are performed individually and independently 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

<b>Alert:</b>	<b>Runtime of the currently running log backups.</b>
Alert ID:	65
Description:	Determines whether or not the most recent log backup terminates in the specified time.
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

<b>Alert:</b>	<b>Storage snapshot is prepared.</b>
Alert ID:	66

<b>Alert:</b>	<b>Storage snapshot is prepared.</b>
Description:	Determines whether or not the period, during which the database is prepared for a storage snapshot, exceeds a given threshold.
Alert Text:	The database was prepared for a storage snapshot for longer than <code>&lt;value&gt; seconds</code> .
User Action:	Investigate why the storage snapshot was not confirmed or abandoned, and resolve the issue.
Default Interval:	5 minutes

<b>Alert:</b>	<b>Enablement of automatic log backup.</b>
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

[Monitoring Alerts \[page 242\]](#)

[Creating Backups \[page 920\]](#)

[Enable or Disable Automatic Log Backup \[page 894\]](#)

## 7.4 Scaling SAP HANA

You can scale SAP HANA either by increasing RAM for a single server, or adding hosts to the system in order to deal with larger workloads. This allows you to go beyond the limits of a single physical server.

There are two general approaches you can take to scale your SAP HANA system.

First, you can scale up. This means increasing the size of one physical machine by increasing the amount of RAM available for processing.

You can also scale out. This 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 means that an SAP HANA system can 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).

The following sections describe the various aspects of scalability.

---

## 7.4.1 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 overall system performance.

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 \[page 372\]](#)

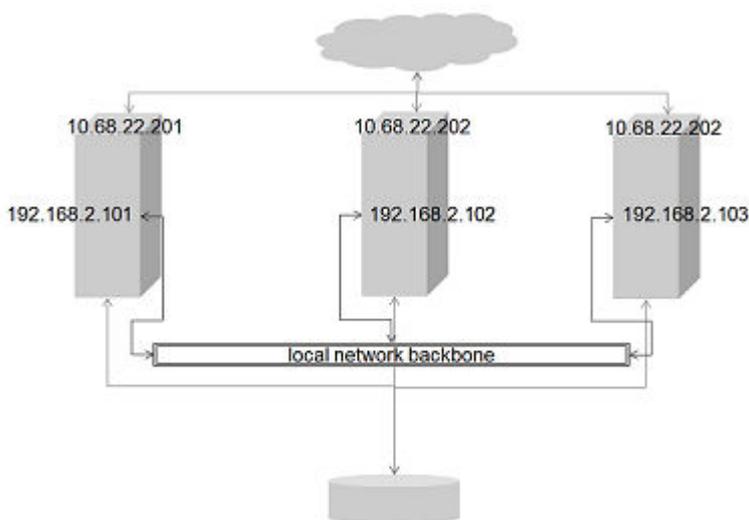
[Workload Management \[page 418\]](#)

## 7.4.2 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 hostnames to IP addresses by editing the section `internal_hostname_resolution` in the `global.ini` file on the [Configuration](#) tab of the SAP HANA studio.

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 hostnames 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.

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.

You will find additional information in the *SAP HANA Security Guide* and the *SAP HANA Master Guide*.

## 7.4.3 Mapping Host Names for Database Client Access

Clients communicate with the database through external hostnames or external IP addresses. A default mapping of external hostnames to internal hostnames 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

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 .`

## 7.4.4 Add or Remove Hosts from an SAP HANA System

You can alter the size of your SAP HANA system by adding or removing hosts.

For detailed information on adding one or more hosts to your SAP HANA System see *Add Hosts to an SAP HANA System*.

For detailed information on removing one or more hosts from your SAP HANA System see *Remove Hosts from an SAP HANA System*.

### Related Information

[Add Hosts Using the Command-Line Interface \[page 541\]](#)

[Remove Hosts Using the Command-Line Interface \[page 549\]](#)

## 7.4.5 Scaling SAP HANA Extended Application Services (XS)

You can configure multiple SAP HANA XS instances to work in a scale out SAP HANA system.

### Context

If you have an application based on SAP HANA Extended Application Services (XS), and you are expecting a high degree of concurrency, you may want to distribute XS 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 \[page 217\]](#)

[Configure HTTP Load Balancing for SAP HANA Extended Application Services \(XS\) \[page 853\]](#)

## 7.4.6 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 &lt;instance_number&gt; -function StartSystem HDB</code>
Stop the system	<code>/usr/sap/hostctrl/exe/sapcontrol -nr &lt;instance_number&gt; -function StopSystem HDB</code>
Query current status of all hosts in the system	<code>/usr/sap/hostctrl/exe/sapcontrol -nr &lt;instance_number&gt; -function GetSystemInstanceList</code>

### **i** Note

HDB start or HDB stop only starts and stops the local host.

## 7.4.7 Table Distribution in SAP HANA

SAP HANA supports different ways of distributing data between multiple index servers in a single system such as database partitioning and table partitioning.

There are two ways that data can be distributed between multiple index servers in a single system:

- Different tables can be assigned to different index servers, which normally run on different hosts (database partitioning).
- A table can be split such that different rows of the table are stored on different index servers (table partitioning).

When a non-partitioned table is created in a distributed system, it must be assigned to one index server. By default, new tables are distributed across available index servers using a round-robin approach. For example, if there are three available index servers A, B, and C (including the master), the first table created will be located on server A, the next one on server B, the next on server C, and so on.

In addition, it is also possible to specify explicitly that a table or a partition be created on a specific index server using the following commands:

- `CREATE TABLE <table_name> <table_contents_source> AT [LOCATION] 'host:port'`
- `CREATE TABLE <table_name> <table_contents_source> <partition_clause> AT [LOCATION] {'host:port' | ('host:port', ...) }`

## Table Distribution for SAP Business Warehouse (BW) Mode

If you are running SAP BW on SAP HANA and you want the table location to be decided by a predefined BW table distribution rule, the `method` parameter in the `table_placement` section of the `global.ini` file must be 2.

All the BW master data tables are distributed across the slave servers using round robin. Fact, DataStore Object (DSO), and Persistent Staging Area (PSA) tables are partitioned with the predefined rule and distributed to the slave servers. Thus, there are no BW data tables on the master server. The other BW tables (all row-store tables, ABAP system tables and general operation tables like REPOSRC ) are created on the master server.

### **i** Note

This parameter is only applied when creating new tables.

This behavior is controlled by two additional parameters.

The parameter `bw_schema` in the `table_placement` section of the `global.ini` controls which schema is allowed on the master server. Tables in all other schemas are created on the slave servers. During installation this parameter is set to the BW schema and the BW shadow schema which is used during migration. Within the BW schema there are ABAP runtime tables and BW data tables. The BW data tables typically start with a namespace. Therefore a slash (“/”) is used to separate these two kinds of data. All tables that start with a slash are always created on the slave servers. All other tables are only created on the master server if they are in one of the BW schemas.

The `prefix` parameter controls the naming convention, which is set to “/” by default.

---

## Using Side-by-Side Replication within an SAP Business Warehouse (BW) Landscape

If you want to run a side-by-side replication scenario with tables replicated from a different system in the landscape there are two options to consider.

1. If the replicated tables are used for BW Virtual InfoProviders, you are advised to treat them like BW tables and locate them on slave servers (this is the default behavior).
2. You might want to have all replicated tables on the master server. In this case, it is required to add the schema name to the parameter `bw_schema` in the `table_placement` section of the `global.ini` file.

For more information, see SAP Note 1637145.

### Related Information

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

[SAP Note 1637145](#)

## 7.4.8 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 (*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 \[page 372\]](#)

[Redistribution of Tables in a Distributed SAP HANA System \[page 1007\]](#)

## 7.4.9 Redistribution of Tables in a Distributed SAP HANA System

In a distributed SAP HANA system, tables and table partitions are assigned to an index server on a particular host at their time of creation, but this assignment can be changed. In certain situations, it is even necessary.

For example, if you plan to remove a host from your system, then you first need to move all the data on that host first to the other hosts in the system. Redistributing tables may also be useful if you suspect that the current distribution is no longer optimal.

Although it is possible to move tables and table partitions manually from one host to another, this is neither practical nor feasible for a large-scale redistribution of data. 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

To plan, adjust and analyze landscape redistribution, you can use the Data Distribution Optimizer. The Data Distribution Optimizer is an SAP HANA XS-based tool included in the SAP HANA Data Warehousing Foundation. 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 *SAP HANA Data Warehousing Foundation - Data Distribution Optimizer Administration Guide*.

## Related Information

[Save Current Table Distribution \[page 1008\]](#)

[Redistribute Tables Before Removing a Host \[page 1009\]](#)

[Redistribute Tables After Adding a Host \[page 1010\]](#)

[Restore Previous Table Distribution \[page 1011\]](#)

[Optimize Table Distribution \[page 1012\]](#)

[Optimize Table Partitioning \[page 1012\]](#)

[Modify Table Distribution Manually \[page 1014\]](#)

[Monitor Table Distribution \[page 1006\]](#)

[Aspects of Scalability \[page 999\]](#)

[Table Placement \[page 414\]](#)

### 7.4.9.1 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.

## 7.4.9.2 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 \[page 920\]](#)

[Add or Remove Hosts from an SAP HANA System \[page 1003\]](#)

## 7.4.9.3 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 \[page 920\]](#)

[Add or Remove Hosts from an SAP HANA System \[page 1003\]](#)

## 7.4.9.4 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.

## 7.4.9.5 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 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 *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.

## 7.4.9.6 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 to 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 \[page 1008\]](#)

## 7.4.9.7 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 \[page 372\]](#)

## 7.4.9.8 Table Redistribution Algorithms

The table redistribution feature uses a number of algorithms to reorganize the placement of the tables and their partitions in the database landscape.

The following algorithms are used.

Table Redistribution Algorithms

ID	Name	Description
1	Add server	After adding one or more index servers to an SAP HANA scale out landscape, partitioned tables should be checked if the changed metric enables additional parts for tables.
2	Clear Server	Prepare removal of servers from landscape
4	Save	Save current landscape setup
5	Restore	Restore a saved landscape setup
6	Balance Landscape/Table	Balances SAP HANA landscapes according to the existing meta information

ID	Name	Description
7	Split	Splits tables if the partitions exceed a configured threshold
8	R3 Load Support	Balances SAP HANA landscapes with externally specified values. Used for migration scenarios or landscape optimization after analyzing application workload.
14	Check Table Placement	Check current landscape against tableplacement rules and (if necessary) provides a plan to correct misalignments.
15	Rerun failed Items	Rerun failed items from previous executed plans

## 8 Maintaining the Application Services Run-Time Environment

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 \[page 1017\]](#)

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1139\]](#)

### 8.1 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 Administration Tools \[page 1018\]](#)

[SAP HANA XS Administration Roles \[page 1020\]](#)

[SAP HANA XS Configuration Parameters \[page 1022\]](#)

### 8.1.1 SAP HANA XS 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"> <li>• Application security (public/private)</li> <li>• User-authentication methods (basic, form-based, logon tickets, X509, SAML)</li> <li>• CORS setup for cross-origin resource sharing</li> <li>• Custom headers: enable support for X-Frame-Options HTTP header</li> <li>• HTTP destinations</li> <li>• SQL connection configurations (for SQL connections for users other than the user specified in the HTTP request).</li> </ul>
<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"> <li>• Management of SAML <b>service</b>-providers, including URLs and metadata management</li> </ul>
<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"> <li>• Management of SAML <b>identity</b>-providers, including IDP metadata, certificates, and destinations</li> </ul>
<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"> <li>• SMTP host settings</li> <li>• Authentication type</li> <li>• Transport security</li> </ul>
<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"> <li>• Trust store configuration and management</li> <li>• Certificate management</li> </ul>
<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"> <li>• Enable the job scheduler</li> <li>• Monitor job-schedule status</li> <li>• Display and maintain schedule's runtime configuration</li> <li>• Add schedules to (or delete from) an XS job</li> </ul>

## Additional SAP HANA XS Tools

The following table lists some tools that are not strictly part of the SAP HANA XS Administration toolset. 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
<i>Online Translation Tool</i>	Maintain translations, for example, for UI text elements	<ul style="list-style-type: none"><li>• Add, modify, delete translation texts</li><li>• Export translation text from SAP HANA to an XML-based xliif-format file</li><li>• Import translation text into SAP HANA.</li></ul>
<i>User Self Service Tools</i>	A set of tools that enable you to maintain user self-service requests and administer the self-service tools themselves.	<ul style="list-style-type: none"><li>• Activate the user self-service tools</li><li>• Maintain user self-service requests</li><li>• Maintain user black/white lists</li><li>• Maintain user self-service e-mail templates</li></ul>

## Related Information

[Maintaining Application Runtime Configurations \[page 1029\]](#)

[Maintaining SAML Providers \[page 1050\]](#)

[Managing Trust Relationships \[page 1043\]](#)

[Maintaining SMTP Server Configurations \[page 1059\]](#)

[Scheduling XS Jobs \[page 1110\]](#)

[Maintaining User Self Service Tools \[page 1085\]](#)

[Maintaining Translation Text Strings \[page 1124\]](#)

## 8.1.2 SAP HANA XS 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

Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. 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 .xsjob 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 .xsjob 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 (.xsoauthclientconfig) 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 <b>outbound</b> communication, for example, the SSL/TLS certificates required by SAP HANA XS applications that connect to an ABAP system.
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 <b>inbound</b> 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.

### 8.1.3 SAP HANA XS Configuration Parameters

An overview of the parameters that the administrator can set to configure how the various components of the XS engine work.

The *xengine.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.

### **i** 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
<a href="#">application_container [page 1023]</a>	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>xscfunc</code> call.
<a href="#">authentication [page 1024]</a>	Options for application-related authentication settings, for example, the location of trust stores.
<a href="#">communication [page 1024]</a>	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
<a href="#">customer_usage [page 1025]</a>	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.
<a href="#">debugger [page 1025]</a>	Settings for the debugging tools, for example, for XS JavaScript.
<a href="#">httpservlet [page 1025]</a>	Options for the SAP HANA XS Web server, for example, port numbers, and maximum number of sessions and threads allowed
<a href="#">odata [page 1026]</a>	Configuration settings for OData requests
<a href="#">scheduler [page 1027]</a>	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

### **i** 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
<a href="#">webdispatcher.ini/profile [page 1027]</a>	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
application_list	Comma-separated list of libraries that can be loaded from an <code>xscfunc</code> call	<code>libxsdxc, InformationAccess</code>	<code>libxsdxc, InformationAccess, libtrustmanager, libxsauthenticator, 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
logonticket_redirect_url	URL that is used to redirect the client to a system that provides SAP logon tickets for SSO authentication	<code>http://link.to.portal/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
default_read_timeout	Time (in milliseconds) before a connection request is closed	-1, 30, 60	-1 (no time set)
default_read_timeout_override	Ignore setting for <code>default_read_timeout</code>	No, Yes	Yes
listenport	The port number on which the XS Web server listens for requests	30007	3\$(SAPSYSTEM)07
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 <b>all</b> 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

Parameter	Description	Example Value	Default Value
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>	<p>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>.</p> <div style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;"> <p><b>→ Tip</b></p> <p>The parameter value <code>icm/HTTP/logging_n</code> is the same as the key defined in the <code>webdispatcher.ini/profile</code> section of the configuration parameters, and "n" must be a unique number.</p> </div>	<code>icm/HTTP/logging_n</code>	<p>N/A</p> <p>The parameter <code>/path/to/XSapp</code> is set (or removed) automatically when the administrator uses the <a href="#">XS Admin Tools</a> to enable (or disable) HTTP tracing on the SAP Web Dispatcher for an application.</p>

## 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"> <li>File must be reachable by http(s)</li> <li>No requirement for authentication or authorization</li> <li>Recommended minimum resolution of image: 1600*1200</li> <li>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 <a href="#">sap.hana.xs.selfService.user.roles.USSExecutor</a>. This user will be used to query the details from the server.</li> </ul>	/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
allow_nullable_keys	Specify if <code>"key"</code> entity elements can ( <code>null</code> ) or cannot ( <code>not null</code> ) have the value NULL.	True/False	False

## 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	300 seconds	900 seconds
<code>disable_job_after_restarts</code>	The maximum number of unsuccessful attempts to start a job before the job schedule is automatically disabled	3	5

## 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"> <li>• PREFIX= the fully qualified path to (and the name of) the application to be traced, for example, sap.hana.xs.admin</li> <li>• LOGFILE= the location of the log file used to store the trace information; the location includes a variable for the application's name (access_log_app-) and the year, month, and day (%y-%m-%d)</li> <li>• MAXSIZEKB= the maximum allowed size and format of the trace file</li> <li>• SWITCHTF=the time of the day when the new log file is created (DAY/NIGHT)</li> <li>• LOGFORMAT= the format of the trace file content, for example: CLF (common log format), CLFMOD (modified CLF), SAP (SAP log file format), SAPSMD, ...</li> <li>• FLUSH=enable or disable the log flush mechanism</li> </ul> <p><b>→ Tip</b></p> <p>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.</p>	<pre>PREFIX=/sap/hana/ide/, LOGFILE=\$( _LOCAL_HOST_NAME) / trace/ access_log_sap.hana.id e=%y-%m-%d, MAXSIZEKB=10000 SWITCHTF=day, LOGFORMAT=SAP, FLUSH=1</pre>	<p>N/A</p> <p>The parameter icm/HTTP/logging_n 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>

## Related Information

<https://service.sap.com/sap/support/notes/1849775>

## 8.1.4 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.

### **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 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 \[page 1030\]](#)

[Edit and SQL Connection Configuration \[page 1035\]](#)

[Edit an HTTP Destination \[page 1037\]](#)

[Maintain XS Job Details \[page 1111\]](#)

### 8.1.4.1 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 \[page 1032\]](#)

[Configure HTTPS \(SSL\) for Client Application Access \[page 1067\]](#)

### 8.1.4.1.1 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 \[page 1033\]](#)
- [CORS \[page 1033\]](#)
- [Custom Headers \[page 1034\]](#)

## 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
<i>Authentication Type</i>	Enables/Disables requirement for user authentication to access an application service. If you <b>enable</b> authentication, you must select the methods that the application applies to authenticate users, for example, SAML or logon tickets.	Public (No Authentication Required)
<i>Connection Security</i>	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
<i>Public Access for Sub-Packages</i>	Enables public access to sub packages in an application package hierarchy. This setting cannot be changed for packages shipped with SAP HANA.	Allowed
<i>Authentication Methods</i>	Defines one or 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
<i>SAML Identity Provider</i>	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
<i>Cross Origin Resource Sharing</i>	Enable/Disable requests from other browser sessions to an application.

CORS Option	Description
<i>ALLOWED ORIGINS</i>	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.
<i>ALLOWED HEADERS</i>	A single header or a comma-separated list of <b>request</b> headers that are allowed by the server. If no request header is specified, no default value is supplied.
<i>EXPOSED HEADERS</i>	A single header or a comma-separated list of <b>response</b> headers that are allowed to be exposed. If no response header is specified for exposure, no default value is supplied.
<i>ALLOWED METHODS</i>	A single permitted method or a comma-separated list of methods that are allowed by the server, for example, <code>GET</code> , <code>POST</code> . If no method is specified, the default <code>GET</code> , <code>POST</code> , <code>HEAD</code> , <code>OPTIONS</code> (all) applies. Note that matching is case-sensitive.
<i>MAX AGE</i>	A single value specifying how long a preflight request should be cached for. If no value is specified, the default time of <code>3600</code> (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
<i>Custom Headers</i>	Enable/Disable the use of custom headers in HTTP response.	Disabled
<i>X-Frame Options</i>	Allow/Deny requests to display content from the same or another Web site. Note that you can only specify <b>one</b> URL with the <code>ALLOW-FROM</code> option, for example: <code>"value": "ALLOW-FROM http://www.site.com".</code>	DENY, SAMEORIGIN, ALLOW-FROM

## Related Information

[Create an Application Runtime Configuration \[page 1030\]](#)

## 8.1.4.2 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 `.xs_sqlcc`. 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 (`XSQLCC_AUTO_USER_...`) for the XSQL 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 \[page 1036\]](#)

### 8.1.4.2.1 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

UI Element	Description	Example
<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 <code>XSSQLCC_AUTO_USER_[...]</code> 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 <b>and</b> ensure that the user has the necessary privileges (for example, by assigning a role).	<code>XSSQLCC_AUTO_USER_[...]</code>
<i>Password</i>	The password for the user bound to the SQL connection configuration. A password is not required for the automatically generated <code>XSSQLCC_AUTO_USER_[...]</code> .	*****
<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
<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>	<code>acme.com.xs.roles::JobAdministrator</code>
<i>Status</i>	The current runtime status of the XSSQL connection configuration (active/inactive)	active

## Related Information

[Edit an SQL Connection Configuration \[page 1035\]](#)

### 8.1.4.3 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/`.

### 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.

### 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 \[page 1040\]](#)

## 8.1.4.3.1 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 \[page 1040\]](#)
- [Proxy Details \[page 1040\]](#)
- [Authentication Details \[page 1041\]](#)
- [OAuth Details \[page 1042\]](#)

### 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

#### **i** 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, Email, Email, 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
<i>SAML Assertion Propagation</i>	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
<i>Entity ID</i>	samlProvider	The entity ID of the remote SAML party
<i>Assertion Consumer Service</i>	samlACS	The way in which SAML assertions or responses are sent
<i>Attributes</i>	samlAttributes	Additional attributes for the SAML Assertion.
<i>User Mapping</i>	samlNameId	A list of name-ID mappings, for example, <i>e-mail</i> .

## 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
<i>OAuth App Config Package</i>	The name of the repository package containing the OAuth application-configuration	sap.hana.xs.oauth.lib.providerconfig
<i>OAuth App Config Name</i>	The name of the OAuth application-configuration (repository artifacts with the suffix <code>.xssoauthappconfig</code> )	abap.xsoauthappconfig

### **i** 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 \[page 1037\]](#)

## 8.1.5 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 (`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 `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 \[page 1045\]](#)

[Import a Server Certificate \[page 1049\]](#)

[Create Your Own Certificate \[page 1046\]](#)

### 8.1.5.1 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 \[page 1049\]](#)

[Create Your Own Certificate \[page 1046\]](#)

[Managing Client Certificates in the SAP HANA Database \[page 758\]](#)

## 8.1.5.2 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/`.

### **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.

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
<b>Certificate collection with purpose <i>SAML</i> (recommended)</b>	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.
<b>Trust store in the file system</b>	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 \[page 1045\]](#)

[Managing Client Certificates in the SAP HANA Database \[page 758\]](#)

## 8.1.5.3 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 \[page 1045\]](#)

[Managing Client Certificates in the SAP HANA Database \[page 758\]](#)

## 8.1.6 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 \[page 1051\]](#)

[Add an SAML Identity Provider \[page 664\]](#)

## 8.1.6.1 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/`.

### 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:

### 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 `<OrganizationDisplay Name>` 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 \[page 1053\]](#)

### 8.1.6.1.1 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 \[page 1053\]](#)
- [Service Provider Configuration \[page 1054\]](#)
- [Metadata \[page 1054\]](#)

#### **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
<i>Name</i>	The fully qualified name of the system hosting the SAML service	SAMLSP01
<i>Organisation Name</i>	The name of the organisation responsible for the SAML service provider. The name you enter is wrapped in the XML tag <code>&lt;OrganizationName&gt;</code> used in the SAML certificate. <i>Organization Name</i> can (but does not have to) be human readable	SAP
<i>Organisation Display Name</i>	The human-readable name of the organisation responsible for the SAML service provider. The name you enter here is wrapped in the XML tag <code>&lt;OrganizationDisplayname&gt;</code> used in the SAML certificate.	SAP
<i>Organisation URL</i>	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
<i>Hash</i>	The hash algorithm use to encode SAML assertions	SHA256
<i>Add Key Info</i>	If <Keyinfo> node should be included in the XML signature; default = yes	"yes" or "no"
<i>Default Application Path</i>	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	/
<i>Assertion Timeout</i>	Period of time (in seconds) for which SAML assertion requests for SSO initiated by an SAML service provider remain valid; default=10 minutes	1000
<i>Default Role</i>	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
<i>Metadata</i>	An XML file containing details of the SAML service provider used to handle SAML assertions

### 8.1.6.2 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/`.

### 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.

### 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 \[page 1051\]](#)

[Modify an Existing SAML Identity Provider \[page 1058\]](#)

[SAML Identity Provider Details \[page 1056\]](#)

[Managing Client Certificates in the SAP HANA Database \[page 758\]](#)

[Configure SSO with SAML Authentication for SAP HANA XS Applications \[page 1078\]](#)

### 8.1.6.2.1 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 \[page 1057\]](#)
- [HTTP Destination \[page 1057\]](#)

## 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.

### 8.1.6.3 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/`.

##### 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 \[page 1054\]](#)

## 8.1.7 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.

### 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 \[page 1062\]](#)

[SAP HANA XS Configuration Parameters \[page 1022\]](#)

### 8.1.7.1 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 \[page 1059\]](#)

[SMTP Configuration Details \[page 1064\]](#)

### 8.1.7.1.1 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 \[page 1064\]](#)
- [Logon authentication type \[page 1064\]](#)
- [Transport security type \[page 1065\]](#)
- [Socket proxy settings \[page 1065\]](#)
- [Other settings \[page 1066\]](#)

## 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.	johndoe
<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 \[page 1062\]](#)

## 8.1.8 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 \[page 1067\]](#)

---

[Maintain Standard HTTP Port Numbers with SAP HANA XS \[page 1069\]](#)

[Configure HTTP\(S\) Access to Multitenant Database Containers \[page 1070\]](#)

## 8.1.8.1 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/`.

#### **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 *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.

## 8.1.8.2 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:

► [SAP HANA Administration Console](#) ► [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](#)

### 8.1.8.3 Configure HTTP(S) Access to Multitenant Database Containers

To enable Web-based applications to send HTTP(S) requests to multitenant database containers, 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 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 in a multiple-container system, 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` or is not available (revisions earlier than SPS 10), 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 (as of SPS 11). 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"> <li>1. Open the Administration editor and choose the <i>Configuration</i> tab.</li> <li>2. Navigate to the <code>xsengine.ini</code> file and expand the <code>public_urls</code> section.</li> <li>3. For each tenant database in the system, add the new properties <code>http_url</code> and <code>https_url</code> at the <b>database layer</b> and enter its public URL as the value: <code>http://&lt;virtual_hostname&gt;:80&lt;instance&gt;</code></li> </ol>
SQL	<p>For each tenant database, execute the statements:</p> <ul style="list-style-type: none"> <li>◦ ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', '&lt;tenant_DB_name&gt;') SET ('public_urls', 'http_url') = 'http://&lt;virtual_hostname&gt;:80&lt;instance&gt;' WITH RECONFIGURE;</li> <li>◦ ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', '&lt;tenant_DB_name&gt;') SET ('public_urls', 'https_url') = 'https://&lt;virtual_hostname&gt;:43&lt;instance&gt;' WITH RECONFIGURE;</li> </ul>

### **i** 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"> <li>1. Start the SAP HANA Web Dispatcher Administration tool at <code>http://&lt;localhost&gt;:80&lt;instance&gt;/sap/hana/xs/wdisp/admin/</code>.</li> <li>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"> <li>1. From the main menu, choose <b>SSL and Trust Configuration</b> <b>&gt;</b> <b>PSE Management</b> <b>▾</b>.</li> </ol> </li> </ol>

Option	Description
	<ol style="list-style-type: none"> <li>2. From the <i>Manage PSE</i> menu, choose <i>SAPSSLS.pse</i>.</li> <li>3. Choose <i>Recreate PSE</i>.</li> <li>4. Enter a distinguished name that matches the host name of all tenant databases. <div data-bbox="485 434 1394 689" style="background-color: #fff9c4; padding: 10px; margin: 10px 0;"> <p> <b>Example</b></p> <ul style="list-style-type: none"> <li>○ Physical host name: myhost.example.com</li> <li>○ Tenant host name 1: mydatabase1.example.com</li> <li>○ Tenant host name 2: mydatabase2.example.com</li> </ul> <p>In this case, you specify <b>CN=*.example.com</b> as the DN, thus creating a server certificate that matches all tenant databases and the system database.</p> </div> </li> <li>5. Choose <i>Create</i>.</li> <li>6. Create a certificate request and submit to your certificate authority (CA) for signing (<i>Create CA Response</i>).</li> <li>3. Import the signed certificate</li> </ol> <p>For more information, see <i>Configure HTTPS (SSL) for Client Application Access</i>.</p>
<b>Individual certificates for each database</b>	<ol style="list-style-type: none"> <li>1. Start the SAP HANA Web Dispatcher Administration tool at <code>http://&lt;localhost&gt;:80&lt;instance&gt;/sap/hana/xs/wdisp/admin/</code>.</li> <li>2. For each tenant database and the system database, create a new trust store with a unique certificate: <ol style="list-style-type: none"> <li>1. From the main menu, choose <b>SSL and Trust Configuration &gt; PSE Management</b>.</li> <li>2. On the PSE management screen, choose <i>Create New PSE</i>.</li> <li>3. Enter a file name for the new PSE. <div data-bbox="485 1122 1394 1234" style="background-color: #fff9c4; padding: 10px; margin: 10px 0;"> <p> <b>Example</b></p> <p>example.pse</p> </div> </li> <li>4. Enter the distinguished name: <b>CN=&lt;host name used for the tenant database in the public_urls section of the xsengine.ini file&gt;</b></li> <li>5. Choose <i>Create</i>.</li> <li>6. For the new PSE, create a certificate request and submit to your CA for signing (<i>Create CA Response</i>).</li> <li>7. Import the signed certificate into the new PSE (<i>Import CA Response</i>).</li> </ol> </li> <li>3. Configure the Web Dispatcher to use multiple certificates: <ol style="list-style-type: none"> <li>1. In the <code>webdispatcher.ini</code> file, create or change the parameter <code>[profile] icm/ssl_config_0</code>, specifying as the value: <b>ID=ssl_config_main, CRED=SAPSSLS.pse, SNI_CREDS=&lt;semicolon (';') separated list of database PSE files&gt;</b></li> <li>2. Add <b>,SSLCONFIG=ssl_config_main</b> to the value of the <code>icm/server_port</code> parameter for the HTTPS port (by default <code>icm/server_port_1</code>).</li> </ol> <div data-bbox="485 1711 1394 1850" style="background-color: #fff9c4; padding: 10px; margin: 10px 0;"> <p> <b>Example</b></p> <pre>icm/server_port_1 = PROT=HTTPS,PORT=4443\$ (SAPSYSTEM),PROCTIMEOUT=600, SSLCONFIG=ssl_config_main</pre> </div> </li> </ol>

## 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 \[page 1067\]](#)

[Using SAP Web Dispatcher for Load Balancing with Tenant Databases \[page 182\]](#)

## 8.1.9 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 in the SAP HANA Database \[page 758\]](#)

[Configure SSO with X.509 Authentication for SAP HANA XS Applications \[page 1076\]](#)

[Configure SSO with SAML Authentication for SAP HANA XS Applications \[page 1078\]](#)

### 8.1.9.1 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.

The SAP HANA cockpit is available at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/admin/cockpit`.

### **i** Note

You must have the privileges required to perform administration tasks with the certificate management apps of 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/`.

### **i** 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 in the SAP HANA Database \[page 758\]](#)

### 8.1.9.2 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.
- If a certificate collection with purpose *SAML* exists, you have authorization to edit it. 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).
- 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.

The SAP HANA cockpit is available at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/admin/cockpit`.

#### **i** Note

You must have the privileges required to perform administration tasks with the certificate management apps of 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 about how to maintain an SAML provider, see *Maintaining SAML Providers*.

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 \[page 1050\]](#)

[Managing Client Certificates in the SAP HANA Database \[page 758\]](#)

## 8.1.9.3 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.

The SAP HANA cockpit is available at the following URL: `http://<WebServerHost>:80<SAPHANAinstance>/sap/hana/admin/cockpit`.

#### **i** Note

You must have the privileges required to perform administration tasks with the certificate management apps of 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>
```

- `<hostname>`  
The hostname of the server issuing/storing the SAP logon tickets
- `<portnumber>`  
The port number accepting connections on the target server issuing/storing the SAP logon tickets
- `</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*.

### **i** 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.

---

## 8.1.9.4 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>")`

## 8.1.10 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 \[page 1086\]](#)

[User Self-Service Roles \[page 1088\]](#)

### 8.1.10.1 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.

### 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.

### 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 \[page 1089\]](#)

[Configure an SMTP Server for User Self-Service Tools \[page 1090\]](#)

[Configure Access to User-Self-Service Administration Tool \[page 1092\]](#)

### 8.1.10.1.1 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

Do not use the repository roles delivered with SAP HANA directly, but instead use them as templates for creating your own roles. 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
sap.hana.xs.selfService.user.roles::USSAdministrator	<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"><li>• System privileges: USER ADMIN</li><li>• Object privileges: SELECT on the tables USERS (SYS) and USER_PARAMETERS (SYS)</li></ul>
sap.hana.xs.selfService.user.roles::USSExecutor	<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>

## 8.1.10.1.2 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 \[page 1086\]](#)

### 8.1.10.1.3 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 \[page 1086\]](#)

[Maintaining SMTP Server Configurations \[page 1059\]](#)

## 8.1.10.1.4 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 \[page 1086\]](#)

[Display all User Self-Service Requests \[page 1097\]](#)

[Maintain User Self-Service Access Lists \[page 1105\]](#)

## 8.1.10.1.5 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 \[page 1095\]](#)

## 8.1.10.1.5.1 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.

### **i** 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
<i>Automatic User Creation</i>	Controls if a user creation request requires approval from user administration. In both cases the administrator has to assign roles to the new user. <ul style="list-style-type: none"> <li><i>Disabled</i>: Requests for a new user account require administrator approval for account activation.</li> <li><i>Enabled</i>: The user is automatically created <b>and activated</b> as a restricted user.</li> </ul>	<code>automatic_user_creation</code>	Disabled/False
<i>Forgot Password</i>	Defines if the system supports password recovery with user-self-service tools. The parameter controls not only the display of the <i>Forgot Password</i> button in the UI logon screen but also the enablement of the corresponding user-self-service backend services.	<code>forgot_password</code>	Disabled/False
<i>Request New user</i>	Enables system support for user-self-service tools. The parameter controls not only the display of the <i>Request New User</i> button in the UI logon screen but also the enablement of the corresponding user-self-service backend services.	<code>request_new_user</code>	Disabled/False
<i>Reset Locked User</i>	Enables support for a password reset for a locked user. Reset password will be forbidden for locked users if the value is <i>Disabled</i> .	<code>reset_locked_user</code>	Disabled/False
<i>Sender E-Mail Address</i>	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.	<code>sender_email</code>	None

UI Element	Description	Parameter Name	Default
<i>Token Expiry Time</i>	The time duration (in seconds) for which a generated token (and the corresponding request for a new user or password reset) is valid.	token_expiry_time	3600
<i>User Creation Request Count</i>	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.	user_creation_request_count	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.xsqlcc`. 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.

### **i** Note

The parameter `login_screen_background_image` must be set in the `httpserver` 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>httpserver</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 \[page 1093\]](#)

## 8.1.10.2 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 `xsengine.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 \[page 1097\]](#)

[Activate a User Account \[page 1102\]](#)

[Reject a User Self-Service Request \[page 1104\]](#)

## 8.1.10.3 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.

## 8.1.10.4 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 \[page 1098\]](#)

### 8.1.10.4.1 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 \[page 1101\]](#)
- [Preferences \[page 1101\]](#)
- [Change Password \[page 1102\]](#)

## 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

UI Element	Description	Example
<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 assignen when the account was activated or, if changed, the currently valid password	*****
<i>New Password</i>	The news password	*****
<i>Repeat Password</i>	Confirm the new password you entered in <i>New Password</i>	*****

## Related Information

[Maintain Your User Profile \[page 1099\]](#)

### 8.1.10.5 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 \[page 1104\]](#)

[Maintain User Self-Service Access Lists \[page 1105\]](#)

[Display all User Self-Service Requests \[page 1097\]](#)

## 8.1.10.6 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 assign 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 \[page 1105\]](#)

[Request a New User Account \[page 1098\]](#)

## 8.1.10.7 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 `xengine.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 \[page 1107\]](#)

## 8.1.10.7.1 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>xsengine.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 <b>restricted</b> 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.
DB Users	The names of the database users who are not allowed to change their respective SAP HANA password using USS reset-password tools, for example, joedoe or janedoe. The following additional restrictions apply: <ul style="list-style-type: none"><li>• By default, it is not possible to use USS tools to reset the password for the SYSTEM user.</li><li>• 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.</li><li>• Users who log on to SAP HANA with SSO credentials cannot use USS tools to reset their password.</li></ul>

## 8.1.10.8 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 \[page 1109\]](#)

### 8.1.10.8.1 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 \[page 1109\]](#)
- [User Activation \[page 1110\]](#)
- [Forgot Password \[page 1110\]](#)

## 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 \[page 1108\]](#)

### 8.1.11 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 \[page 1122\]](#)

[Tutorial: Schedule an XS Job \[page 1119\]](#)

### 8.1.11.1 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**.

Toggleing 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:

- *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 runtime 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).
- 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 \[page 1113\]](#)

[The XS Job File \[page 1122\]](#)

### 8.1.11.1.1 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.

#### **i** Note

toggling the setting for the **Scheduler Enabled** button also changes the current value of the SAP HANA configuration variable `▶ xengine.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).

## Related Information

[Maintain XS Job Details \[page 1111\]](#)

### 8.1.11.1.2 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 \[page 1114\]](#)
- [Runtime Configuration \[page 1115\]](#)
- [Log Cleanup \[page 1116\]](#)

## 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.

### **i** 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.

## **i** 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>XSCron</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
<i>Description</i>	A short description of the XS job schedule	gfn test schedule
<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 xscron 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 xscron job runs	en_US

UI Element	Description	Example
<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 must be assigned the *JobAdministrator* role.

#### XS Job Log Cleanup

UI Element	Description	Example
<i>Enabled</i>	Enable the log-cleanup schedule	Yes
<i>XSCron</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 Administration Roles \[page 1020\]](#)

[Scheduling XS Jobs \[page 1110\]](#)

## 8.1.11.2 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.

## Related Information

[Maintain XS Job Details \[page 1111\]](#)

[The XS Job Dashboard \[page 1113\]](#)

## 8.1.11.3 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`.

#### **i** 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",
      "xscron": "* * * * * 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.

**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`.

You need to specify the following details:

o [User](#)

The user account in which the job runs, for example, **SYSTEM**

o [Password](#)

The password required for user, whose account is used to run the job.

o [Locale](#)

The language encoding required for the locale in which the job runs, for example, **en\_US**

o [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).

o [Active](#)

Enable or disable the job schedule

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 \[page 1122\]](#)

### 8.1.11.3.1 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>",
```

### **i** 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.

### **i** 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 \[page 1119\]](#)

## 8.1.12 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 \[page 1125\]](#)

[Export and Import Translated Text \[page 1129\]](#)

[SAP Translation Hub Cloud Service \(beta\)](#)

### 8.1.12.1 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`.

### 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 provides 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 \[page 1127\]](#)

[Export and Import Translated Text \[page 1129\]](#)

[Edit an HTTP Destination Runtime Configuration \[page 1037\]](#)

[Managing Trust Relationships \[page 1043\]](#)

### 8.1.12.1.1 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.

## 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 Management</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>	<p>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.</p> <div data-bbox="497 1406 1094 1648" style="background-color: #fff9c4; padding: 5px;"><p> <b>Restriction</b></p><p>Access to external translation services is not granted in the SAP HANA license. To use external translation services such as the <i>SAP Translation Hub</i>, an additional license is required. The <i>SAP Translation Hub</i> is currently available only for BETA testing.</p></div> <p>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.</p>	Yes/No

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 \[page 1125\]](#)

[Export and Import Translated Text \[page 1129\]](#)

[Managing Trust Relationships \[page 1043\]](#)

## 8.1.12.2 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`.

### 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 \[page 1127\]](#)

[Create and Edit Text Translations \[page 1125\]](#)

### 8.1.13 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.

#### **i** 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 Administration Roles \[page 1020\]](#)

[Enable HTTP Tracing for an SAP HANA XS Application \[page 1134\]](#)

[Maintain XS Job Details \[page 1111\]](#)

## 8.1.13.1 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/`.

### 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 \[page 1134\]](#)

[Application HTTP Tracing Details \[page 1133\]](#)

### 8.1.13.1.1 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 \[page 1134\]](#)

## 8.1.13.2 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 \[page 1133\]](#)
- [SAP HANA XS Configuration Parameters \[page 1022\]](#)
- [SAP HANA XS Administration Roles \[page 1020\]](#)

## 8.1.13.3 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 Administration Roles \[page 1020\]](#)

[SAP HANA XS Configuration Parameters \[page 1022\]](#)

## 8.2 Maintaining the SAP HANA XS Advanced Model Run Time

Maintain the SAP HANA XS advanced 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 advanced model (XS advanced) run-time environment. In the SAP HANA administration cockpit, the *XS Advanced 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 role "SITE\_DESIGNER". In addition, some of the tools listed below are only available to users to whom the suitable role collection has been assigned. For example, the "XS\_AUTHORIZATION\_ADMIN" role collection includes the authorization scopes required for unrestricted access to the *Application Role Builder* and *SAML Identity Providers Configuration* tools.

- *Application Monitor*  
Monitor the system usage of the applications running in the XS Advanced Model run-time
- *Organization and Space Management*  
Create, list, or delete user organizations and spaces in the XS Advanced Model run time.
- *Application Role Builder*  
Maintain and manage user roles and role collections in SAP HANA.
- *SAML Identity Providers Configuration*  
Configure SAML Identity providers (IDP) for SAP HANA XS advanced model applications that use SAML assertions as the logon authentication method.
- *User Management*  
Create and manage users for SAP HANA XS advanced model applications.
- *SAP HANA Logical Database Setup*  
Manage SAP HANA database instances for SAP HANA XS advanced model applications.
- *SAP HANA Service Broker Configuration*  
Manage and monitor the SAP HANA service broker used by SAP HANA XS advanced model applications.
- *Job Scheduler Service Dashboard*  
Create, schedule, and manage long running operations jobs in the SAP HANA XS advanced model run-time environment.

### **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.

## Related Information

[Monitoring the SAP HANA XS Advanced Model Runtime \[page 1143\]](#)

[Maintaining Organizations and Spaces in SAP HANA XS Advanced Model \[page 1147\]](#)

[Building Roles for SAP HANA XS Advanced Model Applications \[page 1155\]](#)

[Managing SAML Identity Providers in XS Advanced \[page 1162\]](#)

[Scheduling Jobs in XS Advanced \[page 1168\]](#)

## 8.2.1 SAP HANA XS Advanced Administration Tools

SAP HANA XS advanced model includes a Web-based tool that enables you to maintain important parts of the application-development environment, for example, security and authentication methods.

In the administration cockpit, the *XS Advanced Administration* tile catalog includes a selection of Web-based tools that enables you to configure and maintain the basic administration-related elements of the application-development process for the XS advanced run-time environment. The features included in the Web-based SAP HANA *XS Advanced Administration and Monitoring* tool cover the following areas:

Administration Tools for XS Advanced Applications

Tool Name	Description	Scope
<i>Application Monitor</i>	Monitor the system usage of the applications running in the XS advanced model run time.	<ul style="list-style-type: none"> <li>• Management of application security</li> <li>• Monitoring of application resource usage</li> </ul>
<i>Organization and Space Management</i>	Create, list, or delete user organizations and spaces in the XS Advanced Model run time.	<ul style="list-style-type: none"> <li>• Management of organizations</li> <li>• Management of spaces</li> <li>• Management of users in organizations and spaces</li> <li>• Management of XS Advanced business user roles for organizations and spaces</li> </ul>

Tool Name	Description	Scope
<i>SAML Identity Provider</i>	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"> <li>• Management of SAML <b>identity</b>-providers, including IDP metadata, certificates, and destinations</li> </ul>
<i>Application Role Builder</i>	Maintain and manage XS Advanced business user roles and role collections in XS advanced.	<ul style="list-style-type: none"> <li>• Management of XS Advanced business user role collections</li> <li>• Creation of roles and role collections</li> </ul>
<i>User Management</i>	Maintain and manage database users for SAP HANA XS advanced	<ul style="list-style-type: none"> <li>• Creation of XS advanced business users</li> </ul>
<i>SAP HANA Logical Database Setup</i>	Register database instances for use with SAP HANA XS advanced	<ul style="list-style-type: none"> <li>• Registration of databases in XS advanced</li> </ul>
<i>SAP HANA Service Broker Configuration</i>	Manage the SAP HANA service broker in SAP HANA XS advanced	<ul style="list-style-type: none"> <li>• Management of SAP HANA service broker</li> <li>• Creation of roles and role collections</li> </ul>
<i>Job Scheduler Service Dashboard *</i>	Create, schedule, and manage long running operations jobs in the SAP HANA XS advanced model run-time environment.	<ul style="list-style-type: none"> <li>• Enabling of job scheduler services</li> <li>• Monitoring of job-schedule status</li> <li>• Display and maintenance of job schedule's run-time configuration</li> <li>• Creation of schedules for (or deletion from) an XS job</li> </ul>

## Related Information

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1139\]](#)

[Role Collections for XS Advanced Administrators \[page 1141\]](#)

## 8.2.2 Role Collections for XS Advanced Administrators

Access to XS advanced administration tools is granted by means of authorization scopes contained in roles that are grouped into role collections.

To grant a user a specific role in an organisation or space, you can use the `xs` command-line client. For example, to grant a user the role "OrgManager" or "OrgAuditor" in an **organization**, use the `set-org-role` command as admin user in the Admin UI, as illustrated in the following example:

```
xs set-org-role <userName> <orgName> <OrgManager | OrgAuditor>
```

To grant a user the role “SpaceManager” or “SpaceDeveloper” role in a user space, use the `set-space-role` command as admin user, as illustrated in the following example:

```
xs set-space-role <userName> <orgName> <spaceName> <SpaceManager | SpaceDeveloper | SpaceAuditor>
```

### ➔ 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`.

## Access to XS Advanced Administration Tools

In the SAP HANA administration cockpit, the *XS Advanced Administration* tile catalog contains the *Administration and Monitoring* tile, which displays all the tools provided to help you maintain the XS advanced model run-time configurations.

### ⚠ Restriction

In the SAP HANA cockpit, tiles and tile catalogs are only visible to users who have been assigned the role “SITE\_DESIGNER”.

The tools listed in the following table are only available to users to whom the suitable role collection has been assigned; the table shows which role collection is required to use a particular XS-advanced administration tool. For example, the “XS\_AUTHORIZATION\_ADMIN” role collection includes the authorization scopes required for unrestricted access to the *Application Role Builder* and *SAML Identity Providers Configuration* tools.

Roles and Role Collections for XS Advanced Administration

XS Advanced Admin Tool	Role Collection	Comments
<i>Application Monitor</i>	XS_CONTROLLER_ADMIN	Based on the role collection assigned, the user is permitted to perform some, most, or all operations in the application
<i>Organization and Space Management</i>	XS_CONTROLLER_USER or	
<i>SAP HANA Service Broker Configuration</i>	XS_CONTROLLER_AUDITOR	
		<ul style="list-style-type: none"> <li>• ADMIN: No access restrictions</li> <li>• USER: Modify access within the assigned organization or space</li> <li>• AUDITOR: Read-only access within the assigned organization or space</li> </ul>
<i>Role Builder</i>	XS_AUTHORIZATION_ADMIN or XS_AUTHORIZATION_DISPLAY	Based on the role collection assigned the user is permitted to perform some

XS Advanced Admin Tool	Role Collection	Comments
<i>SAML IDP Configuration</i>		(or all) operations in the XS advanced administration tool: <ul style="list-style-type: none"> <li>• ADMIN: Full admin edit access to the tool</li> <li>• DISPLAY Read-only access to the tool</li> </ul>
<i>User Management</i>	XS_USER_ADMIN	Based on the role collection assigned, the user is permitted to perform all operations in the application.
<i>SAP HANA Logical Database Setup</i>	No special roles required.	The user is permitted to perform all operations in the application.

## Related Information

[Maintain Role Collections for XS Advanced Applications \[page 1161\]](#)

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1139\]](#)

## 8.2.3 Monitoring the SAP HANA XS Advanced Model Runtime

Monitor the system usage of the applications running in the XS Advanced Model run-time.

The *Application Monitor* tool enables you to view the system resources used by the individual application instances running in the SAP HANA XS Advanced Model run-time environment. For example, you can see how much memory is allocated and how long the application has been running for.

## Related Information

[Display a List of Running Applications \[page 1144\]](#)

## 8.2.3.1 Display a List of Running Applications

Display a list of all the applications currently active in the SAP HANA XS Advanced Model run time.

### Prerequisites

To access the *Application Monitor* tool, a user requires 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 within an assigned organization or space
- XS\_CONTROLLER\_AUDITOR  
Read-only access within an assigned organization or space

#### → 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

The *Application Monitoring* tool displays detailed information about the resources consumed by all the applications running in the SAP HANA XS Advanced Model run-time. For example, you can see how much memory an application is using, and how this usage changes over time. You can also see how much CPU time is consumed by an application, check where the application is running, and display the port on which the application is reachable.

To display a list of applications running in the selected SAP HANA XS Advanced Model run-time, perform the following steps:

### Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *Application Monitor* tool.
5. Choose the information you want to display. *None* is selected by default and it displays the list all the applications.

6. Display details of an individual application.

In the list of applications, click the application whose details you want to display; the information is displayed in a separate tab.

#### **i** Note

You can choose to display applications based on the Organization assigned to the user, Spaces in the Organization, type of Application and MTA (multi target archive). For this, click *Organization* or *Space* or *Application Type* or *MTA* tab respectively.

7. You can view a graphical representation of all the above information by clicking graph tab.

## Related Information

[Application Monitoring Details \[page 1145\]](#)

[Monitoring the SAP HANA XS Advanced Model Runtime \[page 1143\]](#)

[Role Collections for XS Advanced Administrators \[page 1141\]](#)

### 8.2.3.1.1 Application Monitoring Details

The information available in the *Applications* tab in the *Application Monitor* tool for XS advanced administrators.

The *Applications* tab in the *Application Monitor* tool enables you to view details of the XS applications running in the XS advanced run time that you are are monitoring. The following table indicates which information can be viewed.

#### **i** Note

You can choose what information to display and sort the displayed information according to any of the displayed details, for example: *Memory (KB)* usage or *CPU Time*.

XS Application Details

UI Element	Description	Example
<i>Application</i>	The name of the application running in the XS advanced run time	jobscheduler-db
<i>Memory (KB)</i>	The amount of memory currently in use by the selected application	1,448
<i>Delta Memory (KB)</i>	The difference between the amount of memory currently in use by the selected application and the amount of memory used by the application the last time the memory usage was measured	0

UI Element	Description	Example
<i>CPU Time (ms)</i>	The amount of CPU time (in milliseconds) used by the application	00:01:13
<i>Delta CPU Time (ms)</i>	The difference (in milliseconds) between the amount of CPU time currently consumed by the selected application and the amount of CPU time used by the application the last time the CPU time was measured	0
<i>User Mode Time (ms)</i>	The amount of user-mode time (in milliseconds) consumed by the 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.	70,526
<i>Kernel Mode Time (ms)</i>	The amount of kernel-mode time (in milliseconds) consumed by the application. In kernel mode, an application has unrestricted access to CPU instructions and memory addresses.	2,800
<i>URL</i>	The URL used to start the listed application. Additional credentials might be required to authenticate the user who is requesting access to the listed application, for example, a user name and the corresponding password.	https://host-name.acme.com:40304
<i>Host *</i>	The name of the host where the listed instance of an XS advanced application (or service) is running	hostname.acme.com
<i>Port *</i>	The port number on which the application (or service) instance is reachable	40304
<i>Instance *</i>	The number of application instances currently running. For more information about a particular application instance, choose [>] at the right-hand end of the instance row.	1, 2, ...
<i>Process ID **</i>	The ID of the process(es) started by the selected application	2456
<i>Statistic Time **</i>	The time at which the most recent measurement was taken	Dec 8, 2015, 2:23:21 PM
<i>Process Priority **</i>	<p>The priority assigned to the application's process(es)</p> <ul style="list-style-type: none"> <li>Linux Lower numbers represent a higher priority; a high priority means more favorable scheduling. Negative numbers (for example, -1) represent real-time priorities, as displayed with the <code>ps</code> (and <code>top</code>) command and the priority option, for example: <code>ps -o priority</code></li> <li>Windows Higher numbers represent a higher priority. If no priority is defined for a process, a value of 0 should be used.</li> </ul>	20

UI Element	Description	Example
<i>Command</i> **	<p>The command used to start the selected process</p> <p>→ <b>Tip</b></p> <p>Click the link to display the full command (including the path to the executable) in a pop-up window.</p>	<code>/usr/sap/hana/ start_node.sh</code>
<i>Parent ID</i> **	The ID of the parent process (if any such process exists) which spawned the listed application's processes	1302

### Restriction

\* This information is only available at the **instance** level.

\*\* This information is only available at the **process** level.

## Related Information

[Display a List of Running Applications \[page 1144\]](#)

## 8.2.4 Maintaining Organizations and Spaces in SAP HANA XS Advanced Model

Create, list, or delete user organizations and spaces in the SAP HANA XS Advanced Model run time.

Organizations enable developers to collaborate by sharing resources, services, and applications. Access to the shared resources, services, and applications is controlled by roles, for example, "Org Manager" or "Org Auditor"; the role defines the scope of the permissions assigned to the named user in the organization. For example, an Org Manager can add new users to organizations; create, modify, or delete organizational spaces; and add domains to the organization.

In an organization, spaces enable users to access shared resources that can be used to develop, deploy, and maintain applications. Access to the resources is controlled by roles, for example, "Space Manager", "Space Developer" or "Space Auditor"; the role defines the scope of the permissions assigned to the named user in the organizational space. For example, a Space Developer can deploy and start an application.

You can use the *Organizations* administration tool to perform the following tasks in the SAP HANA XS Advanced Model run time:

- List, create, rename, and delete organizations
- List, create, rename, and delete organizational spaces
- Maintain user accounts in organizations and spaces

## Related Information

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1139\]](#)

### 8.2.4.1 Maintain Organizations

Create, list, or delete user organizations in the SAP HANA XS Advanced Model run time.

#### Prerequisites

To access the *Organization and Space Management* tool, a user requires 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 within an assigned organization or space
- XS\_CONTROLLER\_AUDITOR  
Read-only access within an assigned organization or space

#### → 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

Organizations enable developers to collaborate by sharing resources, services, and applications. Access to the shared resources, services, and applications is controlled by roles, for example, “Org Manager” or “Org Auditor”; the role defines the scope of the permissions assigned to the named user in the organization. For example, an Org Manager can create and manage users, create, modify, or delete organizational spaces, and add domains to the organization.

To add an organization, perform the following steps:

#### Procedure

1. Start the SAP HANA cockpit.

2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *Organization and Space Management* tool.
5. Add a new organization.

At the bottom of the *Organizations* pane, click the **[+]** button.

6. Define the name of the new organization.

In the *Create Organization* dialog, type a name for the new organization, for example, **myOrg** and choose **OK**.

## Related Information

[Maintaining Organizations and Spaces in SAP HANA XS Advanced Model \[page 1147\]](#)

[Maintain Organizational Spaces \[page 1149\]](#)

[Role Collections for XS Advanced Administrators \[page 1141\]](#)

### 8.2.4.2 Maintain Organizational Spaces

Create, list, or delete user spaces in the SAP HANA XS Advanced Model run time.

## Prerequisites

To access the *Organization and Space Management* tool, a user requires 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 within an assigned organization or space
- XS\_CONTROLLER\_AUDITOR  
Read-only access within an assigned organization or space

#### ➔ 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

In an organization, you define spaces to provide shared resources that can be used to develop, deploy, and maintain applications. Access to the resources is controlled by roles, for example, "Space Manager", "Space Developer" or "Space Auditor". The role defines the scope of the permissions assigned to the named user in the organizational space. For example, a Space Developer can deploy and start an application.

To add a space to an organization, perform the following steps:

## Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *Organization and Space Management* tool.
5. Select the organization to which you want to add a space.

In the *Organizations* list, choose the organization to which you want to add a new space.

6. Add a new space to your organization.
  - a. In the *Organization* details pane, choose *Edit*.
  - b. Choose *[+ Create Space]*.
  - c. In the *Create Space* dialog, type a name of the new space, for example, **mySpace** and choose *OK*.
7. Save your changes.

In the *Organization* pane, choose *Save*.

## Related Information

[Maintaining Organizations and Spaces in SAP HANA XS Advanced Model \[page 1147\]](#)

[Maintain Organizations \[page 1148\]](#)

[Role Collections for XS Advanced Administrators \[page 1141\]](#)

## 8.2.4.3 Maintain Users in Organizations and Spaces

Add new users to organizations and assign user roles.

### Prerequisites

To access the *Organization and Space Management* tool, a user requires 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 within an assigned organization or space
- XS\_CONTROLLER\_AUDITOR  
Read-only access within an assigned organization or space

#### → 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

Roles are used to control user access to the shared resources, services, and applications available in an organization or space, for example, “Org Manager” or “Org Auditor”. The role defines the scope of the permissions assigned to the named user in the selected organization. For example, an Org Manager can create and manage users, create, modify, or delete organizational spaces, and add domains to the organization.

To add a user to an organization, perform the following steps:

### Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *Organization and Space Management* tool.
5. Select the organization where you want to maintain users.

In the *Organizations* list, choose the organization where you want to maintain users. When you select an organization from the list, the users already configured for the selected organization are displayed in the

---

*Users* tab in the *Organization* pane. The *Users* icon also indicates how many users are configured in the selected organization.

6. Display a list of users in the selected organization.

Choose *Users* to display a list of all SAP HANA users configured in the currently selected organization.

7. Add a new user.

- a. In the *Users* details pane, choose *Edit*.
- b. Choose *[+ Add User]*.
- c. In the *Add Users* dialog, search for and select one or more users to add and choose *Add*.

8. Assign a role to the new organization user.

You can choose between the following organization user roles:

- "OrgManager"
- "OrgAuditor"

9. Save your changes.

In the *Organization* pane, choose *Save*.

New users are displayed in the list of users under the role assigned to them.

10. Maintain users in an organizational **space**.

- a. In the *Organizations* list, select the organization containing the space whose users you want to maintain.
- b. In the *Organization* pane, choose the space whose users you want to maintain.
- c. Add a new user to the space.

Choose *Edit* and *[+] Add User*, and in the *Add Users* list select one or more users and choose *Add*.

The new users are displayed in the *Users* pane. At this point the users do not have any roles assigned.

- d. Assign a role to a user.

In the *Users* pane, choose the role (or roles) you want to assign to the new user, for example:

- "SpaceManager"
- "SpaceDeveloper"
- "SpaceAuditor"

11. Save your changes.

In the *Space* pane, choose *Save*.

New users are displayed in the list of users together with the role (or roles) assigned to them.

## Related Information

[Organization and Space Users \[page 1153\]](#)

[Maintaining Organizations and Spaces in SAP HANA XS Advanced Model \[page 1147\]](#)

[Maintain Organizations \[page 1148\]](#)

[Maintain Organizational Spaces \[page 1149\]](#)

[Role Collections for XS Advanced Administrators \[page 1141\]](#)

## 8.2.4.3.1 Organization and Space Users

A list and description of typical users for organizations and spaces in XS advanced.

You can grant or restrict access to organizations and spaces in XS advanced by assigning roles in specific organizations and spaces. The following table lists the roles that you can assign to XS advanced users in a specified organization:

Organizational Roles in XS Advanced

Role	Description
<i>OrgManager</i>	<ul style="list-style-type: none"><li>• Create and manage organization users</li><li>• Create, modify, or delete organizational spaces</li><li>• Add domains to the organization</li></ul>
<i>OrgAuditor</i>	<ul style="list-style-type: none"><li>• View all users in the organization</li><li>• View the roles assigned to a user (or users) in the organization</li><li>• View quotas configured for the organization</li></ul>

The following table lists the roles that you can assign to XS advanced users in a specified space:

Space Roles in XS Advanced

Role	Description
<i>SpaceManager</i>	<ul style="list-style-type: none"><li>• Manage users in the selected space</li><li>• View details of applications running in the space (for example: status, instances, service bindings, and resource usage)</li></ul>
<i>SpaceDeveloper</i>	<ul style="list-style-type: none"><li>• Deploy, start, stop an application</li><li>• Bind an application to (or unbind an application from) a service</li><li>• View details of applications running in the space (for example: status, instances, service bindings, and resource usage)</li></ul>
<i>SpaceAuditor</i>	<ul style="list-style-type: none"><li>• View details of applications running in the space (for example: status, instances, service bindings, and resource usage)</li></ul>

### Related Information

[Maintain Users in Organizations and Spaces \[page 1151\]](#)

[Maintaining Users in XS Advanced \[page 1190\]](#)

## 8.2.5 Setting Up Security Artifacts

Developers create authorization information for business users in their environment; the information is deployed in an application and made available to administrators who complete the authorization setup and assign the 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 in a dedicated XS advanced application. The XS advanced administrators view the authorization information in role templates, which they use as part of the run-time configuration. The administrators use the role templates to build roles, which are aggregated in role collections. The role collections are assigned, in turn, to business users.

The tasks required to set up authorization artifacts in SAP HANA XS advanced are performed by two distinct user roles: the application developer and the SAP HANA administrator. After the deployment of the authorization artifacts as role templates, the administrator of the SAP HANA XS advanced application uses the role templates provided by the developers for building role collections and assigning them to business users in the *SAP HANA XS Administration Tools* section of the *SAP HANA Administration Guide*.

### **i** Note

To test authorization artifacts after deployment, developers can use the role templates to build role collections and assign authorization to business users in the [SAP HANA XS Administration Tools](#).

#### Setting Up Authorization Artifacts

Step	Task	User Role	Tool
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 <code>xsuaa</code> service in XS advanced using the service broker	Application developer	XS advanced CLI tool
4	Bind the service instance to the XS advanced application by including it into 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 in the XS advanced application role builder using role templates	XS advanced administrator	Application role builder
7	Create a role collection and assign roles to it	XS advanced administrator	Application role builder
8	Assign the role collection to a SAML 2.0 identity provider or to SAP HANA database users	XS advanced administrator	Application role builder, and SAML IDP Tool
9	Assign the users to roles using the role collections	XS advanced administrator	User interface of XS Advanced

## Related Information

[Building Roles for SAP HANA XS Advanced Model Applications \[page 1155\]](#)

[Maintain Role Collections for XS Advanced Applications \[page 1161\]](#)

[Add a Role Collection to an SAML IDP \[page 1167\]](#)

[Assign Roles to a User \[page 717\]](#)

## 8.2.6 Building Roles for SAP HANA XS Advanced Model Applications

Maintain application roles and role collections which can be used in user management.

You can use the XS advanced model administration tools to create and manage user roles. The user roles are derived from role templates that are defined in the security description (`xs-security.json`) of applications that have been registered as OAuth 2.0 clients at 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. The roles you create with the XS advanced administration tools can be added to role collections, which can then be assigned to SAP HANA database users or users logging on with SAML 2.0 assertions.

With the *Application Role Builder* tool, you can perform the following tasks:

- Create and delete user roles  
User roles define authorization scopes and are based on the role templates and scopes defined in the application's security description file
- Add user roles to one or more "role collections"
- Configure and manage role collections

### → Tip

The role collections you configure can be assigned to SAP HANA database users or to users logging on with SAML 2.0 assertions.

## Related Information

[Maintain Roles for XS Advanced Applications \[page 1156\]](#)

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1139\]](#)

## 8.2.6.1 Maintain Roles for XS Advanced Applications

Roles are used to define the type of access granted to an application.

### Prerequisites

To access the *Application Role Builder* tool, a user requires the authorization scopes defined in the roles grouped together in one of the following role collections:

- XS\_AUTHORIZATION\_ADMIN  
Full access: no access restrictions
- XS\_AUTHORIZATION\_DISPLAY  
Read-only access to the *Application Role Builder* tool

#### → 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

A role is an instance of a role template; you can build 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. 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. Attributes define information that comes with the respective user, for example 'cost center' or 'country'. This information can only be resolved at run time.

### Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *Application Role Builder* tool.
5. Select the application for which you want to create a role.

In the *Applications* list, choose the application to which you want to add a new role.

The *Application* pane displays information about the selected application in four tabs:

#### i Note

The information displayed is defined in the application's security description file (`xs-security.json`).

- *Roles*  
The roles that have been built

#### **i** Note

Roles are derived from role templates defined in the corresponding application's security description (`xs-security.json`)

- *Scopes*  
A list of the authorization scopes defined in the application's `xs-security.json` file
- *Attributes*  
A list of the attributes defined in the application's `xs-security.json` file and that are applied to the selected role
- *Templates*  
A list of the role templates defined in the application's `xs-security.json` file

#### 6. Create a new role.

- In the *Application* details pane, choose a list of the role templates defined in the application's *Roles*.
- Choose *Create Role*.
- Type a name for the new role.
- Select a role template (from the drop-down list) to use to create the new role.

The list of the role templates is derived from the selected application's `xs-security.json` file.

- Type a short description of the new role.
- Choose a source type for any attributes that need to be applied to the role.

Attributes are defined in the corresponding application's `xs-security.json` file. You can choose between the following attribute source *type*:

- *Static*  
Enter an attribute value manually.
- *SAML Attribute*  
The value of the attribute must be obtained from a SAML 2.0 token.

- Save your changes.

#### 7. Add the new role to a role **collection**. (optional)

You can either add the role to an **existing** role collection or create a **new** role collection to which you add the new role, too.

To add the new role to an existing role collection, perform the following steps:

- Display the list of roles available for the selected application.

In the *Application* pane, choose *Roles*.

- Select the role that you want to add to a role collection.

In the *Role Name* list, you can choose one or more roles to add to one (or multiple) role collection(s).

- Add the selected role(s) to a role collection.

Choose *Add to Role Collection*. In the *Select Role Collections* dialog, select one or more role collections to which you want add the new role, and choose *OK*.

### → Tip

To add a role to a **new** role collection, use the *Configure Role Collections* tool; in the *Role Collections* pane, choose [\[+\]](#) to add a new role collection; choose *Edit*, and [\[+\] Add Application Role](#) to select the application role(s) you want to add, and save the changes.

8. Confirm that the new role has been successfully added to the role collection.
  - a. Switch to the *Role Collection* tool.
  - b. In the *Role Collections* pane, select the role collection that you want to check.
  - c. In the *Role Collections* details pane, confirm that the role you added appears in the list of roles assigned to the role collection.

## Related Information

[Application Role Builder Details \[page 1158\]](#)

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1139\]](#)

[Role Collections for XS Advanced Administrators \[page 1141\]](#)

### 8.2.6.1.1 Application Role Builder Details

The *Application Role Builder* tool for XS advanced model run-time administrators displays detailed information concerning application roles and access.

The *Application Role Builder* includes the following tools:

- *Application Role Builder*  
A role is an instance of a role template; you can build a role based on a role template and assign the role to a role collection.
- *Role Collection*  
Roles are assigned to role collections which are assigned, in turn, to SAP HANA users or SAML2 groups.

## Application Role Builder

A role is an instance of a role template; you can build a role based on a role template and assign the role to a role collection. The *Application Role Builder* tool displays information about the selected application and any related roles in the following windows, tabs, and panes:

- *Roles*
- *Scopes*
- *Attributes*
- *Templates*

## **i** Note

The information displayed is defined in the application's corresponding security descriptor (`xs-security.json`).

## Roles

The *Roles* tab in the *Application* pane provides the following details concerning the application currently selected in the *Applications* list.

Application Role Details

Field Name	Description
<i>Role Name</i>	The name of the role built from a role template
<i>Role Template</i>	The name of the template used to build the selected role
<i>Description</i>	A short explanation of the limits of the selected role
<i>Add to Role Collection</i>	Start the <i>Role Collection</i> wizard and add the selected role(s) to a role collection
<i>Create Role</i>	Start the <i>Add Role</i> wizard and create a new role based on the role templates and authorization scopes defined in the corresponding application's <code>xs-security.json</code> file
<i>Configure Role Collections</i>	Open the <i>Role Collection</i> pane where you can edit the selected role collection, for example, to edit or remove roles from the collection

## Scopes

The *Scopes* tab in the *Application* pane provides the following details concerning the user authorization scopes defined in the security descriptor (`xs-security.json`) of the application currently selected in the *Applications* list.

Application Role Authorization Scope Details

Field Name	Description
<i>Scope</i>	The name of the authorization scope built into a role
<i>Description</i>	A short explanation of the authorization scope defined in the selected role
<i>Configure Role Collections</i>	Open the <i>Role Collection</i> pane where you can edit the selected role collection, for example, to edit or remove roles from the collection

## Attributes

The *Attributes* tab in the *Application* pane provides the following details concerning the attributes to apply to the user roles built from user templates that are defined for the application currently selected in the *Applications* list. 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".

#### Application Role Attributes Details

Field Name	Description
<a href="#">Attribute</a>	The name of the attribute defined to provide details in a role
<a href="#">Description</a>	A short explanation of the authorization scope of the selected role
<a href="#">Configure Role Collections</a>	Open the <a href="#">Role Collection</a> pane where you can edit the selected role collection, for example, to edit or remove roles from the collection

## Templates

The [Templates](#) tab in the [Application](#) pane provides the following details concerning the role templates defined in the security descriptor (`xs-security.json`) of the application currently selected in the [Applications](#) list.

#### Application Role Template Details

Field Name	Description
<a href="#">Role Template Name</a>	The name of the template used to build the selected role
<a href="#">Scope References</a>	A list of the scope references in the role templates defined in the selected application's security descriptor ( <code>xs-security.json</code> )
<a href="#">Attribute References</a>	A list of the attributes that are referenced in the role template defined in the selected application's security descriptor ( <code>xs-security.json</code> )

## Role Collection

Roles are assigned to role collections which are assigned in turn to SAP HANA users or SAML2 groups. The [Role Collection](#) tool displays information about the role collections that have been maintained as well as the roles available in a role collection. Additional information includes: which templates the roles are based on, and which applications the roles apply to. Role collections enable you to group together the roles you create with the XS advanced administration tools; the role collections you define can be assigned to SAP HANA database users or users logging on with SAML assertions.

#### Role Collection Details

Field Name	Description
<a href="#">Application Name</a>	The name of the application to which the selected role is assigned
<a href="#">Role Template</a>	The name of template used to build the role
<a href="#">Role Name</a>	The name of the role in the role collection
<a href="#">Configure Application Roles</a>	Start the <a href="#">Application Role</a> tool and display the list of roles and role templates associated with an application

## Related Information

[Maintain Roles for XS Advanced Applications \[page 1156\]](#)

### 8.2.6.2 Maintain Role Collections for XS Advanced Applications

Role collections group together different roles that can be applied to the application users.

#### Prerequisites

To access the *Application Role Builder* tool, a user requires the authorization scopes defined in the roles grouped together in one of the following role collections:

- XS\_AUTHORIZATION\_ADMIN  
Full access: no access restrictions
- XS\_AUTHORIZATION\_DISPLAY  
Read-only access to the *Application Role Builder* tool

#### → 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

A role is an instance of a role template; you can build 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.

#### Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *Application Role Builder* tool.
5. Display the *Role Collections* tool.

Choose *Configure Role Collection* in the application details pane.

6. Select the role collection that you want to maintain.

Either choose an existing role collection from the list displayed in the *Role Collections* pane, or choose [\[+\]](#) to create a new role collection.

7. Add an application role to the role collection.
  - a. In the *Role Collection* pane, choose *Edit*.
  - b. Choose [\[+\] Add Application Role](#).
  - c. In the *Add Application Role* dialog, select an **application name** from the drop-down list.
  - d. Select a **template** from the drop-down list.
  - e. Select an **application role** from the drop-down list.
  - f. Choose *OK* to save the details.

The change or update is displayed in the *Role Collection* pane.

## Related Information

[Building Roles for SAP HANA XS Advanced Model Applications \[page 1155\]](#)

[Managing SAML Identity Providers in XS Advanced \[page 1162\]](#)

[Role Collections for XS Advanced Administrators \[page 1141\]](#)

## 8.2.7 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 administration tools for *SAP HANA XS advanced model* includes the *SAML Identity Provider Configuration* tool, which you can use to configure SAML Identity providers (IDP) for SAP HANA XS advanced model run time. 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 use the *SAML Identity Provider Configuration* tool to perform the following tasks:

- Add a new SAML Identity provider (IDP)
- Modify the details of an existing SAML Identity provider (IDP)
- Manage role collections based on SAML assertions

### **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

[Maintain an SAML IDP in XS Advanced \[page 1163\]](#)

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1139\]](#)

## 8.2.7.1 Maintain an SAML IDP in XS Advanced

An SAML identity provider (IDP) is used by the SAML service provider to authenticate users signing in to applications by means of SSO.

### Prerequisites

To access the *SAML Identity Provider Configuration* tool, a user requires the authorization scopes defined in the roles grouped together in one of the following role collections:

- XS\_AUTHORIZATION\_ADMIN  
Full access: no access restrictions
- XS\_AUTHORIZATION\_DISPLAY  
Read-only access *SAML Identity Provider Configuration* tool

#### → 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

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.

### Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *SAML Identity Provider Configuration* tool.
5. 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

advanced to use. The only information you must provide manually is the “Origin Key” of the new IDP. Note that the *Entity ID* 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 *Origin Key*, *Subject*, *Entity ID*, and *Issuer* fields in the *General* screen area, and the URL fields in the *Destination* screen area, for example, *Base URL* and *SingleSignOn URL (\*)*.

- c. In the *Origin Key* box of the *General* screen area, check that the Origin Key for the new SAML SSO identity provider has been inserted automatically.

#### **i** Note

The Origin Key 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 advanced applications to use.

6. Save the updated details of the SAML identity provider.

The new SAML IDP is displayed in the *SAML Identity Provider List* pane.

## Related Information

[SAML Identity Providers Details in XS Advanced \[page 1164\]](#)

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1139\]](#)

[Role Collections for XS Advanced Administrators \[page 1141\]](#)

### 8.2.7.1.1 SAML Identity Providers Details in XS Advanced

An SAML identity provider (IDP) is used by the SAML service provider (SP) 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](#)
- [Role Collections](#)

## General

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>Metadata</i>	The text of the SAML certificate	<?xml version="1.0" encoding="utf-8"?> <ns3:EntityDescriptor xmlns:
<i>Origin Key</i>	The name (Origin Key) of the SAML identity provider is mandatory and must be unique.	ACCOUNTS_ACME_COM
<i>State</i>	The current state of the selected SAML Identity Provider (IDP)	Active/Inactive
<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>Name</i>	The name (ID) of the remote SAML party	accounts.acme.com
<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

## Role Collections

The *Role Collections* tab in the *SAML Identity Provider* displays details of role collections configured for an application. Role collections group together different roles that can be assigned to SAP HANA database users or to users logging on with SAML assertions.

## SAML Assertion Role Collection Details

UI Component	Description
<i>Role Collection</i>	<p>The names of the assertion-based role collections associated with the selected application.</p> <p>➔ <b>Tip</b></p> <p>To change the role collection, choose <i>Edit</i> and select the new role collection from the drop-down list.</p>
<i>Attributes</i>	<p>A list of the attributes defined in the selected application's security configuration (<code>xs-security.json</code>) file.</p> <p>⚠ <b>Restriction</b></p> <p>Currently, the only attribute allowed is "Groups".</p>
<i>Operator</i>	<p>The operator to apply in the rule in which the attribute is used.</p> <p>⚠ <b>Restriction</b></p> <p>Currently, the only operator allowed is "equals".</p>
<i>Value</i>	<p>The value of the attribute to use for the rule that triggers the assignment of the selected role collection.</p> <p>➔ <b>Tip</b></p> <p>To change the attribute value, choose <i>Edit</i> and type the new value in the <i>Value</i> box.</p>

## Related Information

[Maintain an SAML IDP in XS Advanced \[page 1163\]](#)

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1139\]](#)

[Maintain Role Collections for XS Advanced Applications \[page 1161\]](#)

## 8.2.7.2 Add a Role Collection to an SAML IDP

Roles grouped in a role collection can be added to an SAML IDP; the collections can be assigned to users on logon.

### Prerequisites

To access the *SAML Identity Provider Configuration* tool, a user requires the authorization scopes defined in the roles grouped together in one of the following role collections:

- XS\_AUTHORIZATION\_ADMIN  
Full access: no access restrictions
- XS\_AUTHORIZATION\_DISPLAY  
Read-only access the *SAML Identity Provider Configuration* tool

#### → 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

You can assign roles to users who log on to an application by means of single sign-on (SSO) with SAML assertions. The roles are defined in one or more role collections, which you can configure and maintain with the *Application Role Builder* tool. The role collections can be added to an SAML IDP.

### Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *SAML Identity Provider Configuration* tool.
5. Choose the SAML identity provider (IDP) to which you want to assign a role collection.  
In the *SAML Identity Provider List*, choose the SAML IDP to which you want to add a role collection.
6. Assign a role collection to the SAML IDP.
  - a. In the *SAML Identity Provider* pane, display the *Role Collections* tab.
  - b. Select the role collection to add.

Choose *Edit* and *[+] Add*, and in the *Add Role Collection* column, use the drop-down list to select the role collection to add to the SAML IDP.

→ Tip

The role collections displayed are the same ones you maintain in the *Application Role Builder* tool.

- c. Define a value to use for the rule to trigger that assignment of the role collection.
- 7. Save the updated details of the SAML identity provider.

Choose *Save* to display details of the updated SAML IDP in the *SAML Identity Provider List* pane.

## Related Information

[Maintain an SAML IDP in XS Advanced \[page 1163\]](#)

[SAML Identity Providers Details in XS Advanced \[page 1164\]](#)

[Maintaining the SAP HANA XS Advanced Model Run Time \[page 1139\]](#)

[Role Collections for XS Advanced Administrators \[page 1141\]](#)

## 8.2.8 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.

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
3	Bind an Application to the Job Scheduler Service	Space Developer
4	Maintain jobs and job schedules	Administrator

## 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

The 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*.

### **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 \[page 1170\]](#)

[Job Scheduler REST API for XS Advanced \[page 1173\]](#)

[The Job Scheduler Dashboard \[page 1189\]](#)

### 8.2.8.1 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 \[page 1173\]](#)

[The Job Scheduler Dashboard \[page 1189\]](#)

[Scheduling Jobs in XS Advanced \[page 1168\]](#)

### 8.2.8.1.1 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
<a href="#">Job Creation</a>	Used to create a job. Job creation can accept a collection of job schedules to be created.	admin
<a href="#">Job Configuration</a>	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.	admin
<a href="#">Job Deletion</a>	Delete a job and purge all its run time information such as job schedules and logs.	admin
<a href="#">Job Schedule Creation</a>	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.	admin
<a href="#">Job Schedule Modification</a>	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.	admin
<a href="#">Job Schedule Deletion</a>	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.	admin
<a href="#">Bulk Job Schedule Activation</a> <a href="#">Bulk Job Schedule Deactivation</a>	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
<a href="#">Job Details</a>	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
<a href="#">Job Schedule Details</a>	Retrieve the saved details and configurations of a specified job schedule & optionally the generated logs for the schedule.	view
<a href="#">Bulk Job Schedule Deletion</a>	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
<a href="#">Job Run Log Update</a>	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

```
POST /scheduler/jobs HTTP/1.1
content-type:application/json;charset=utf-8
host: https://scheduler.service.acme.com
content-length: 500
{"name":"hello_world", "description": "greetts the world periodically",
"action":"http://httpbin.org/basic-auth/abc/def", "active": true,
"httpMethod":"GET", "schedules": [{"repeatEvery":"2 minutes", "description":
"every 2 minutes, run this schedule", "data":{"time":"abc"}, "active": true},
{"cron":"* * * * *", "description": "every minute, run this schedule", "data":
{"time":"abc"}, "active": true}]}
```

```
Response:
Status: 201 CREATED
Location: /scheduler/jobs/109
Content-Type: application/json; charset=utf-8
Body: {"_id":109,"name":"hello_world","description":"greetts the world
periodically","action":"http://httpbin.org/basic-auth/abc/
def","active":true,"user":null,"httpMethod":"GET","schedules":
[{"scheduleId":"a66cbdd4-42ce-4c36-b61a-66bc8be6c2d0","description":"every 2
minutes, run this schedule","data":
{"time":"abc"},"type":"recurring","active":true,"startTime":null,"endTime":null
,"repeatInterval":"2 minutes"}, {"scheduleId":"24d6f8f0-d156-4d48-
b678-44d353e700d2","description":"every minute, run this schedule","data":
{"time":"abc"},"type":"recurring","active":true,"startTime":null,"endTime":null
,"cron":"* * * * *"}]}
```

The job schedule creation request is defined with the parameters listed in the following table:

### 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

**Note**

If a job with the same name for the technical user credentials already exists, the job creation request fails.

Request Field	Type	Description
description	String	Describes the user-defined job
action *	String	The fully qualified URL endpoint to be called when the job runs, for example: <code>http://host.acme.com/app/call</code>
active	Boolean	Defines if the job should be activated on creation. Allowed values are: <ul style="list-style-type: none"> <li>• <code>false</code> (default) The job is in inactive mode on creation</li> <li>• <code>true</code> The job is activated on creation</li> </ul>
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.
repeatEvery	string	For recurring schedules, the parameter denoting when the schedule should run. The parameter supports the use of human readable formats.

Schedule Field	Type	Description
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"
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
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":null
, "repeatInterval":"2 minutes"}, {"scheduleId":"2f98471c-26de-4293-ae53-
e4a16e1513f5", "description":"Default Schedule", "data":
{"time":"abc"}, "type":"recurring", "active":false, "startTime":null, "endTime":nul
l, "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"> <li>false (default) The job is in inactive mode when configured</li> <li>true The job is active when configured</li> </ul>
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.

Request Field	Type	Description
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

```
POST /schedule/jobs/:jobId/schedules HTTP/1.1
content-type:application/json;charset=utf-8
host:https://scheduler.service.acme.com
content-length: 500
```

```
{ "repeatEvery": "2 hours", "data": { "param": "abc" }, "active": true,
  "description": "New Schedule", "startTime": { "date": "2015-04-21", "format":
  "YYYY-MM-DD" } }
```

```
Response:
Status: 200 OK
Content-Type: application/json; charset=utf-8
{ "scheduleId": "0252c255-8f57-4fd5-aa47-c7b344ac9a46", "name": "greet the
world2", "data": { "param": "abc" }, "type": "recurring", "priority":
0, "action": "http://httpbin.org/basic-auth/abc/
def", "nextRunAt": "2015-04-23T03:04:32.856Z", "startTime": "2015-04-20T18:30:00.00
0Z", "endTime": null, "repeatInterval": "2 hours", "active": true, "description": "New
Schedule", "jobId": "132" }
```

#### 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.
repeatEvery	string	For recurring schedules, the parameter denoting 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"> <li>false (default) The job is in inactive mode when configured</li> <li>true The job is active when configured</li> </ul>
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":"2013-02-08T04:00:26.12
3Z","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.
repeatEvery	string	For recurring schedules, the parameter denoting 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
cron	string	For recurring schedules, the parameter denoting the cron pattern. It must be a valid crontab format, for example: " <code>* * * * * */10</code> "
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: <code>{ "dataParam": "somevalue" }</code>
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"> <li><code>false</code> (default) The job is in inactive mode when configured</li> <li><code>true</code> The job is active when configured</li> </ul>
description	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"><li>• false (default) All job schedules for the specified job should be <b>de</b>activated</li><li>• true All job schedules for the specified job should be activated</li></ul>

## 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,"schedu
leId":"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"><li>• false Do <b>not</b> display details of job schedules for the specified job</li><li>• true Display details of job schedules for the specified job</li></ul>
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,"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":nul
l,"endTime":null,"repeatAt":null,"logs":
[{"text":null,"httpStatus":null,"executionTime":null,"status":"SCHEDULED","sche
duleTime":"2015-04-19T15:10:53.000Z","completionTime":null}], "scheduleId":"550d
1b96-8002-4d0d-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`

### 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:

### 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:	<p>"10 hours from now"</p> <p>"20 minutes from now"</p> <p>"in 2 hours"</p> <p>"tomorrow at 4pm"</p> <p>"next week monday at 5am"</p> <p>"9pm tonight"</p> <p>"3.30pm"</p>
<code>repeatAt</code>	<p>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.</p> <div style="background-color: #fff9c4; padding: 10px; margin-top: 10px;"> <p><b>i Note</b></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>	<p>"4.40pm"</p> <p>"18.40"</p> <p>"6.20am"</p> <p>"17.20:30 "</p> <p>"09:30:26.123+07:00"</p>
<code>repeatEvery</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".	<p>"10 hours "</p> <p>"2 days "</p> <p>"3 seconds"</p>

## 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
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)
ggggg	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

Input Token	Example	Description
SS	0-99	Hundredths of a second
SSS	0-999	Thousandths of a second
ZZZ	+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)

## 8.2.8.1.2 The Job Scheduler Dashboard

The Job Scheduler dashboard provides schedule snapshots, view service instances, and manage jobs for a service instance.

You access the Job Scheduler dashboard from the *XS Advanced Administrator and Monitoring Tools* in the SAP HANA administration cockpit. You need specific Job Scheduler roles to perform the Job Scheduler tasks. 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*.

The Job Scheduler dashboard contains the following screens:

- [Home](#)  
 This screen enables you to view and access job scheduler instances in a Space within an Organization.
- [Dashboard](#)  
 This screen provides a snapshot of all the schedules assigned to a particular service instance. This is a read-only view. It presents a graphical representation of the status or the state of the schedule. You can view schedules and their logs. It also provides advanced search capabilities. You also have the option to search for schedules within a specific duration.
- [Manage](#)  
 This screen enables you to manage jobs for the service instance that you selected from the home screen. You can view logs and previous runs of schedules. You can also add or delete schedules.
- [Configuration](#)  
 This screen enables you to maintain global configuration required for functioning of the service.

#### ➔ Tip

To access the *Job Scheduler Dashboard*, you must have the authorization scopes defined in the roles grouped together in one of the default role collections, for example, *XS\_CONTROLLER\_ADMIN*, or

`XS_CONTROLLER_USER`. 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`.

You also have the option to save the service instance that you manage and monitor frequently as a preference.

## 8.2.9 Maintaining Users in XS Advanced

Create, set up and maintain database users in XS advanced.

The XS advanced administrator tools include a *User Management* tool that enables you to perform the following tasks:

- Create a new database user
- Delete an existing database user
- Modify a database user
- Search for a database user by name (or part thereof)

### Related Information

[Create a New User for XS Advanced \[page 1190\]](#)

[Delete an Existing XS Advanced Business User \[page 1194\]](#)

[Modify an XS Advanced Business User Access \[page 1195\]](#)

[Search for an XS Advanced Business User \[page 1196\]](#)

[User Details in XS Advanced \[page 1192\]](#)

### 8.2.9.1 Create a New User for XS Advanced

Create a new database user in SAP HANA XS advanced.

#### Prerequisites

To create a new business user in XS advanced, you require the authorization scope defined in the following role collection:

- `XS_USER_ADMIN`  
Full access: no access restrictions

---

## Context

You can create new XS advanced business users by using the *User Management* tool.

To create a new business user in XS advanced, perform the following steps:

## Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *User Management* tool.
5. Choose *New* to display the *New User* dialog.
6. Enter the details for the new user and choose *Create*.

### **i** Note

An asterisk (\*) indicates a compulsory field, where you must provide information.

## Results

A new business user is created in XS advanced along with a corresponding SAP HANA user.

### **i** Note

The new SAP HANA user is a **restricted** database user with a role collection assigned that only enables access to XS advanced.

## Related Information

[User Details in XS Advanced \[page 1192\]](#)

[Enable XS Advanced Access to Existing SAP HANA User \[page 1192\]](#)

## 8.2.9.1.1 User Details in XS Advanced

Define details of the user to create in XS advanced.

You can display details of an XS advanced business user with the XS advanced administration tools. The details displayed in the *User Management* tool are the same as the information you provided when using the *New User* dialog to create a new business user.

Create New User Details

UI Element	Description
<i>User Name</i>	Name of the XS advanced user
<i>First Name</i>	First name of the user
<i>Last Name</i>	Last name of the user
<i>Email ID</i>	The e-mail address of the user to whom the account and profile belong. E-mail address must be unique to the user
<i>Password</i>	Type a password for the new user's account.; the password must contain a minimum of eight alphanumeric characters
<i>Confirm Password</i>	Type the password again for confirmation

### Related Information

[Create a New User for XS Advanced \[page 1190\]](#)

[Maintaining Users in XS Advanced \[page 1190\]](#)

## 8.2.9.2 Enable XS Advanced Access to Existing SAP HANA User

Enable access to XS advanced for an existing SAP HANA user.

### Prerequisites

To enable XS advanced access for an existing SAP HANA user, you must have the following permissions assigned:

- XS\_USER\_ADMIN  
Full access: no access restrictions
- SAP HANA user

---

The user who requires access to XS advanced must already have a user account in the SAP HANA database.

## Context

An existing SAP HANA user can be given access to XS Advanced. For that follow below steps:

## Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *User Management* tool.
5. Choose *New* to display.
6. In the *New User* dialog, choose the option *Based on Existing HANA User*.
7. Select an SAP HANA user from the list of users displayed.
8. Choose *Create*.

## Results

The selected SAP HANA user is granted access to XS advanced.

## Related Information

[User Details in XS Advanced \[page 1192\]](#)

[Maintaining Users in XS Advanced \[page 1190\]](#)

---

## 8.2.9.3 Delete an Existing XS Advanced Business User

Remove an existing XS advanced business user.

### Prerequisites

To delete an existing user in XS advanced, you require the authorization scopes defined in the following role collection:

- XS\_USER\_ADMIN  
Full access: no access restrictions

### Context

You can delete an XS advanced business user by using the *User Management* tool, as described in the following steps:

### Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *User Management* tool.
5. Choose the check box corresponding to the name of the user that you intend to delete.
6. Choose *Delete*.
7. Confirm that you want to delete the selected user.

### Results

The specified XS advanced business user is deleted.

### Related Information

[User Details in XS Advanced \[page 1192\]](#)

[Maintaining Users in XS Advanced \[page 1190\]](#)

## 8.2.9.4 Modify an XS Advanced Business User Access

Modify details of the access permissions granted to an SAP HANA XS advanced business user.

### Prerequisites

To modify an existing user in XS advanced, you require the authorization scopes defined in the following role collection:

- XS\_USER\_ADMIN  
Full access: no access restrictions

### Context

You can modify details of a user's access to XS advanced with the *User Management* tool, as described in the following steps:

#### **i** Note

It is not possible to use the *User Management* to modify user name.

### Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *User Management* tool.
5. Select the user that you intend to modify.
6. Make the desired changes to the user's access permissions.

You can modify the role collections assigned to the user. To modify, select *Role Collections*. Select *Add* to add new roles and *Remove* to remove existing roles.

#### **i** Note

The role collection *XS\_USER\_PUBLIC* is assigned by default to all XS advanced business users; it and cannot be deleted.

7. Save the changes you made.

---

## Results

You have successfully modified user access.

## Related Information

[User Details in XS Advanced \[page 1192\]](#)

[Maintaining Users in XS Advanced \[page 1190\]](#)

## 8.2.9.5 Search for an XS Advanced Business User

Search for a user in SAP HANA XS advanced.

### Prerequisites

To search for an XS advanced user, you require the authorization scopes defined in the following role collections:

- XS\_USER\_ADMIN  
Full access: no access restrictions
- XS\_USER\_DISPLAY  
Read only access

### Context

To search for a user with the *User Management* tool, perform the following steps:

### Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *User Management* tool.
5. You can search for a user by user name, email ID, first or last name, or assigned role collections.

---

## Results

The system displays the user you are searching for or indicates that the user is not registered.

## Related Information

[User Details in XS Advanced \[page 1192\]](#)

[Maintaining Users in XS Advanced \[page 1190\]](#)

## 8.2.10 Maintaining Database Instances in XS Advanced

Maintain logical database instances in XS advanced.

The XS advanced administrator tools include a *SAP HANA Logical Database Setup* tool, which you can use to perform the following tasks:

- Search for a logical database by name (or part of a name) in XS advanced
- Display the status of all logical database currently available in XS advanced
- Prepare a logical database for use with XS advanced

## Related Information

[Prepare a Logical Database for use with XS Advanced \[page 1197\]](#)

[Search for a Logical Database in XS Advanced \[page 1199\]](#)

### 8.2.10.1 Prepare a Logical Database for use with XS Advanced

Step-by-step instructions for preparing a logical database for use in SAP HANA XS advanced.

## Context

To make a logical database available for use in XS advanced, you must prepare the logical database, for example, using the *SAP HANA Logical Database Setup* tool, as described in the following steps:

## Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *SAP HANA Logical Database Setup* tool.
5. Locate the logical database that you want to “prepare” for use in XS advanced.
6. In the *Prepare* column, choose *Prepare*.

The *System Credential* dialog appears requesting logon credentials for the database user “SYSTEM”.

7. Enter a password for the database SYSTEM user.
8. During the preparation of the logical database, the status *In Progress* is displayed. When the logical database is available for use in XS advanced, the status changes to *Prepared*.

## Results

The selected logical database is “prepared” available for use in XS advanced.

## Related Information

[Database Configuration Details in XS Advanced \[page 1198\]](#)

### 8.2.10.1.1 Database Configuration Details in XS Advanced

A list and description of the details required to configure a logical database in XS advanced.

You can use *SAP HANA Logical Database Setup* tool to prepare a new logical database for use with XS advanced. The following table indicates the information displayed for the logical database

Logical Database Details

UI Element	Description	Example
<i>Name</i>	The name of the logical database.	<b>TenantDB1</b> <b>MyDatabaseName</b>
<i>ID</i>	The ID of the prepared logical database	

UI Element	Description	Example
<i>Prepare</i>	The status of the logical database, for example, <i>Prepared</i> .	<ul style="list-style-type: none"> <li>• <i>Prepared</i> The selected logical database is available for use with XS advanced</li> <li>• <i>Prepare</i> The selected logical database is <b>not yet</b> ready for use with XS advanced; choose <i>Prepare</i> to start the preparation process</li> <li>• <i>In Progress</i> Preparation of the logical database for use with XS advanced is not yet complete</li> </ul>

## Related Information

[Maintaining Database Instances in XS Advanced \[page 1197\]](#)

[Prepare a Logical Database for use with XS Advanced \[page 1197\]](#)

[Search for a Logical Database in XS Advanced \[page 1199\]](#)

## 8.2.10.2 Search for a Logical Database in XS Advanced

List the logical databases in XS advanced whose name matches a specified search string.

### Context

You can use the *Search* tool to find an existing logical database by name (or any part of the name). If you enter only part of a name, the *Search* tool filters the list of logical databases and displays only those whose names include the string you type. For example, if you type **DB** in the *Search* tool, the list of databases displayed in *Logical Databases (#)* is restricted to any names that contain "DB", for example: "MyTenantDB", "DB1", or "MyLogicalDB5".

To use the *Search* tool to filter the list of logical database displayed according to a specified string, perform the following steps:

### Procedure

1. Start the SAP HANA cockpit.

2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring* tools.
4. Start the *SAP HANA Logical Database Setup* tool.
5. In the *Search* field enter the name of the logical database you want to display, for example, **MyLogicalDB** or any part of the name, for example, **DB**.

## Results

The *Search* tool displays a list of logical databases that match the string you specified.

## Related Information

[Database Configuration Details in XS Advanced \[page 1198\]](#)

[Prepare a Logical Database for use with XS Advanced \[page 1197\]](#)

## 8.2.11 Maintaining the Service Broker in XS Advanced

Map a logical database to an organization or space in XS advanced.

The XS advanced administrator tools include a *SAP HANA Service Broker Configuration* tool that enables you to map a logical database to an organization or space in XS advanced. All applications are deployed to the logical database in the mapped organization or space.

## Related Information

[Configure the Service Broker in XS Advanced \[page 1200\]](#)

[Service Broker Configuration Details in XS Advanced \[page 1201\]](#)

### 8.2.11.1 Configure the Service Broker in XS Advanced

Map a logical database to an organization or space in XS Advanced.

## Context

The *SAP HANA Service Broker Configuration* tool enables you to map a logical database (for example, created using the *SAP HANA Logical Database Setup* tool) to a particular organization or a space within the specified

---

organization. When the mapping is completed, all the applications in the mapped organization or space are deployed to the target logical database.

### **i** Note

Only one logical database can be mapped to an organization or space.

To map a logical database to an organization or space in XS advanced, perform the following steps:

## Procedure

1. Start the SAP HANA cockpit.
2. Locate the *XS Advanced Administration* tools catalog.
3. Start the *Administration and Monitoring Tools*.
4. Start the *SAP HANA Service Broker Configuration* tool.
5. In the list of *Organizations and Spaces*, select an organization or space.
6. In *Logical Database*, choose a logical database from the list displayed to map it to the selected organization or space.

## Results

If you map a logical database at the organization level, the *Logical Database* field for the spaces in the specified organization are disabled. You can change the mapping at any time by choosing *Reset*.

## Related Information

[Service Broker Configuration Details in XS Advanced \[page 1201\]](#)

[Prepare a Logical Database for use with XS Advanced \[page 1197\]](#)

### 8.2.11.1.1 Service Broker Configuration Details in XS Advanced

Enter details of the SAP HANA Service Broker Configuration in XS Advanced.

You can map a logical database to an Organization or Space in XS advanced with the help of *SAP HANA Service Broker Configuration* tool. The mapping ensures that all applications in the mapped organization or space in organization are deployed to the corresponding logical database.

### Restriction

A logical database can be mapped to only one organization or space.

The *Logical Database Mapping* tool in *SAP HANA Service Broker Configuration* requires the following details:

Service Broker Configuration Details

UI Element	Description	Example
<i>Organization and Spaces</i>	A list of the organizations and spaces currently configured and available in XS advanced	<b>orgname/PROD</b> <b>orgname/MySpace</b>
<i>Logical Databases</i>	A list of all the logical database currently configured and available in XS advanced. Use the <i>Search Logical Database</i> dialog to locate and select the logical database you want to map to the organization or space indicated in <i>Organizations and Spaces</i> .	<b>MyDatabaseName</b>

## Related Information

[Configure the Service Broker in XS Advanced \[page 1200\]](#)

[Maintaining Database Instances in XS Advanced \[page 1197\]](#)

## 8.2.12 Maintaining the SAP HANA Deployment Infrastructure (HDI) Configuration

Set up and manage the SAP HANA Deployment Infrastructure (HDI).

You can use parameters to control the execution flow of SAP HANA DI (HDI) procedure calls and, in addition, manage the behavior of various parts of the SAP HANA Deployment Infrastructure, for example, HDI containers. This section provides information about the following parameters:

- **SAP HANA DI Parameters**  
Control the execution flow of SAP HANA DI procedures and SAP HANA DI container-specific procedures. You can also configure parameters that control the execution flow of the deployment process of a build plug-in for all database objects of the corresponding type.
- **SAP HANA DI Configuration Parameters**  
Configure the general behavior of SAP HANA DI and, in addition, the behavior of HDI containers.

## Related Information

[SAP HANA DI Parameters \[page 1203\]](#)

[SAP HANA DI Configuration Parameters \[page 1213\]](#)

### 8.2.12.1 SAP HANA DI Parameters

Overview of available SAP HANA DI and build plugin parameters.

In SAP HANA DI, parameters are a means of controlling the execution flow of SAP HANA DI procedure calls. There are three types of parameters in SAP HANA DI:

- **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 plug-in 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.
- **Path parameters**  
A path parameter determines the execution flow of the deployment process of a single database artifact. For example, a path parameter can be used to specify the batch size for batched database access or for batch processing of a **specific** database artifact.

## 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
<code>_SYS_DI.CONFIGURE_CONTAINER_PARAMETERS</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.CONFIGURE_DI_PARAMETERS</code>	trace_context trace_level.<trace topic>

SAP HANA DI Call	Available Parameters
<code>_SYS_DI.CONFIGURE_LIBRARIES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> undeploy
<code>_SYS_DI.CREATE_CONTAINER</code>	trace_context trace_level.<trace topic>
<code>_SYS_DI.DROP_CONTAINER</code>	container_lock_wait_timeout ignore_deployed ignore_errors ignore_work trace_context trace_level.<trace topic>
<code>_SYS_DI.EXPORT_CONTAINER</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.GRANT_CONTAINER_API_PRIVILEGES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.GRANT_CONTAINER_API_PRIVILEGES_WITH_GRANT_OPTION</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.GRANT_CONTAINER_SCHEMA_PRIVILEGES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.GRANT_CONTAINER_SCHEMA_ROLES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.GRANT_CONTAINER_SUPPORT_PRIVILEGE</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.IMPORT_CONTAINER</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.LIST_CONFIGURED_LIBRARIES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.LIST_LIBRARIES</code>	trace_context trace_level.<trace topic>

SAP HANA DI Call	Available Parameters
<code>_SYS_DI.REVOKE_CONTAINER_API_PRIVILEGES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.REVOKE_CONTAINER_SCHEMA_PRIVILEGES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.REVOKE_CONTAINER_SCHEMA_ROLES</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.REVOKE_CONTAINER_SUPPORT_PRIVILEGE</code>	container_lock_wait_timeout trace_context trace_level.<trace topic>
<code>_SYS_DI.CONFIGURE_CONTAINER</code>	Deprecated since SAP HANA SPS 12.
<code>_SYS_DI.CONFIGURE_DI</code>	Deprecated since SAP HANA 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>&lt;container&gt;#DI.CANCEL</code>	container_lock_wait_timeout trace_context trace_level.<trace topic> treat_warnings_as_errors

SAP HANA DI Container-Specific Call	Available Parameters
<container>#DI.CONFIGURE_CONTAINER_PARAMETERS	container_lock_wait_timeout trace_context trace_level.<trace topic>
<container>#DI.CONFIGURE_LIBRARIES	container_lock_wait_timeout trace_context trace_level.<trace topic> undeploy
<container>#DI.DELETE	container_lock_wait_timeout ignore_non_existing_paths recursive trace_context trace_level.<trace topic>
<container>#DI.GRANT_CONTAINER_API_PRIVILEGES	container_lock_wait_timeout trace_context trace_level.<trace topic>
<container>#DI.GRANT_CONTAINER_API_PRIVILEGES_WITH_GRANT_OPTION	container_lock_wait_timeout trace_context trace_level.<trace topic>
<container>#DI.GRANT_CONTAINER_SCHEMA_PRIVILEGES	container_lock_wait_timeout trace_context trace_level.<trace topic>
<container>#DI.GRANT_CONTAINER_SCHEMA_ROLES	container_lock_wait_timeout trace_context trace_level.<trace topic>
<container>#DI.LIST	container_lock_wait_timeout ignore_files ignore_folders recursive trace_context trace_level.<trace topic>
<container>#DI.LIST_CONFIGURED_LIBRARIES	container_lock_wait_timeout trace_context trace_level.<trace topic>
<container>#DI.LIST_DEPLOYED	container_lock_wait_timeout ignore_files ignore_folders recursive trace_context trace_level.<trace topic>

SAP HANA DI Container-Specific Call	Available Parameters
<container>#DI.MAKE	simulate_make trace_context trace_level.<trace topic> treat_warnings_as_errors
<container>#DI.MAKE_ASYNC	simulate_make trace_context trace_level.<trace topic> treat_warnings_as_errors
<container>#DI.READ	container_lock_wait_timeout ignore_files ignore_folders recursive trace_context trace_level.<trace topic>
<container>#DI.READ_DEPLOYED	container_lock_wait_timeout ignore_files ignore_folders recursive trace_context trace_level.<trace topic>
<container>#DI.REVOKE_CONTAINER_API_PRIVILEGES	container_lock_wait_timeout trace_context trace_level.<trace topic>
<container>#DI.REVOKE_CONTAINER_SCHEMA_PRIVILEGES	container_lock_wait_timeout trace_context trace_level.<trace topic>
<container>#DI.REVOKE_GRANT_CONTAINER_SCHEMA_ROLES	container_lock_wait_timeout trace_context trace_level.<trace topic>
<container>#DI.STATUS	container_lock_wait_timeout trace_context trace_level.<trace topic>
<container>#DI.WRITE	container_lock_wait_timeout trace_context trace_level.<trace topic>
<container>#DI.CONFIGURE_CONTAINER	Deprecated since SAP HANA 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_timeout	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> .
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.

Parameter	Possible values	Description
<code>ignore_errors</code>	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.
<code>ignore_files</code>	true, false	Indicates if files are to be ignored in the output when reading files.  The default value is false.
<code>ignore_folders</code>	true, false	Indicates if folders are to be ignored in the output when reading files.  The default value is false.
<code>ignore_non_existing_paths</code>	true, false	Indicates if paths that do not exist should be ignored, for example, when deleting folders.  The default value is false.
<code>ignore_work</code>	true, false	Indicates if existing files in the work file system are to be ignored when dropping a container.  The default value is false.
<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 only be simulated.  The default value is false.

Parameter	Possible values	Description
<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_&lt;trace topic&gt;</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.&lt;trace topic&gt;</code>	Fatal, Error, Warning, Info, Interface, Debug, InterfaceFull, DebugFull	Specifies the trace level of a specific trace topic. <trace topic> may be an arbitrary SAP HANA trace topic.  There is no default value.
<code>treat_warnings_as_errors</code>	true, false	Indicates if warnings during an SAP HANA DI call should be treated as errors.  The default value is false.
<code>undeploy</code>	true, false	Indicates if, in case of a call to configure libraries, files corresponding to a library to be removed should also be undeployed.  The default value is false.

## Parameters for Build Plugins

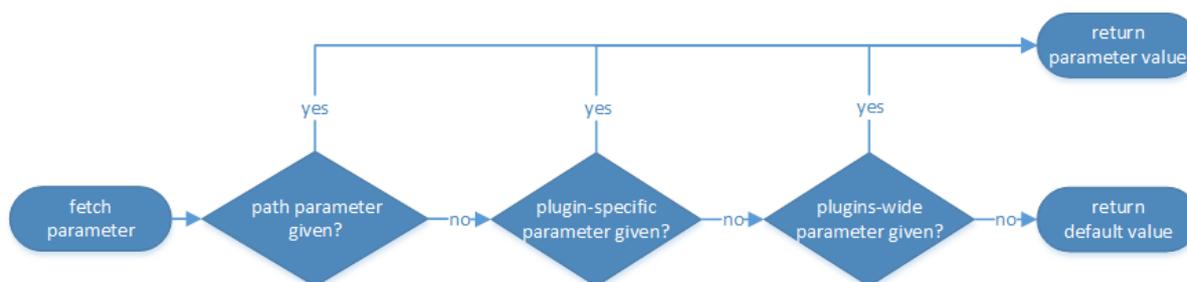
SAP HANA DI supports three types of parameters for controlling the execution flow of build plugins. On a global level, a "plugins-wide parameter" applies to all build plugins supporting the parameter. On a more fine-grained level, a "plugin-specific parameter" only applies to the specified build plugin. 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	-	<code>com.sap.hana.di/&lt;key&gt;</code>
plugin-specific parameter	-	<code>com.sap.hana.di.&lt;plugin&gt;/&lt;key&gt;</code>

Parameter Type	File	Parameter Structure
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 build plugin parameters.

## Build Plugins

The following table lists the available parameters for the build plugins.

Build Plugin	Available Build Plugin Parameters	Available Path Parameters
com.sap.hana.di.cds	force_undeploy	-
com.sap.hana.di.role	force_undeploy	-

Build Plugin	Available Build Plugin Parameters	Available Path Parameters
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.
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.

## Related Information

[SAP HANA DI Configuration Parameters \[page 1213\]](#)

### 8.2.12.2 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
api.severity_for_invalid_parameter	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.

Parameter	Possible Values	Description
<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.
<code>blobs.transaction_lock_wait_timeout</code>	0 ... 2,147,483,647	Specifies the time (in milliseconds) a blob store operation waits for a locking conflict to clear.  The default value is 100,000.
<code>connection.container_default_transaction_lock_wait_timeout</code>	0 ... 2,147,483,647	Specifies the default 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.global_transaction_lock_wait_timeout</code> .
<code>connection.global_transaction_lock_wait_timeout</code>	0 ... 2,147,483,647	Specifies the time (in milliseconds) an SAP HANA DI operation waits for a locking conflict to clear.  The default value is 100,000.
<code>connection.max_polls_for_master_indexserver</code>	0 ... 2,147,483,647	Specifies the maximum number of polls for the master indexserver before aborting the SAP HANA DI operation.  The default value is 100.
<code>connection.poll_interval_for_master_indexserver</code>	0 ... 2,147,483,647	Specifies the interval (in seconds) between polls for the master indexserver.  The default value is 5.
<code>make.default_max_parallel_jobs</code>	0 ... 2,147,483,647	Specifies the default maximum number of parallel jobs to be spawned during a make.  The default value is 16.

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>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
blobs.days_to_keep	-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 blobs.container_default_days_to_keep.
connection.transaction_lock_wait_timeout	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 connection.container_default_transaction_lock_wait_timeout.

Parameter	Possible Values	Description
make.max_parallel_jobs	0 ... 2,147,483,647	Specifies the maximum number of parallel jobs to be spawned during a make.  The default value is the value of the corresponding SAP HANA DI configuration parameter make.default_max_parallel_jobs.
messages.days_to_keep	-2,147,483,648 ... 2,147,483,647	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.  The default value is the value of the corresponding SAP HANA DI configuration parameter messages.container_default_days_to_keep.
messages.requests_to_keep	-2,147,483,648 ... 2,147,483,647	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.  The default value is the value of the corresponding SAP HANA DI configuration parameter messages.container_default_requests_to_keep.

## 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, ?, ?, ?);
```

---

## 9 SAP HANA Data Provisioning

SAP HANA integrates remote data from many data sources to enrich your applications and to deliver in-depth analysis.

This section covers the following data acquisition features in SAP HANA:

- **Smart Data Access (SDA)**  
SAP HANA smart data access 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.
- **Hadoop Integration**  
You can also use SAP HANA to access data from Apache Hadoop.
- **Smart Data Integration**  
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.

### Related Information

[SAP HANA Smart Data Access \[page 1218\]](#)

[SAP HANA Hadoop Integration \[page 1258\]](#)

[Data Replication and Transformation \[page 1268\]](#)

### 9.1 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 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.

As part of the HANA core system, no additional licensing is required to use smart data access. However, additional installation packages must be downloaded and installed using the SAP HANA database lifecycle manager (HDBLCM).

The following remote data sources are supported:

- SAP HANA

- SAP IQ
- SAP Adaptive Server Enterprise
- SAP Event Stream Processor (supported on Intel-based hardware platforms only)
- SAP MaxDB (supported on Intel-based hardware platforms only)
- Apache Hadoop (supported on Intel-based hardware platforms only)
- Teradata Database (supported on Intel-based hardware platforms only)
- Microsoft SQL Server 2012 (supported on Intel-based hardware platforms only)
- Oracle Database 12C
- IBM DB2 (supported on Intel-based hardware platforms only)
- IBM Netezza Appliance (supported on Intel-based hardware platforms only)
- Apache Spark

Get started with remote data sources as follows:

1. Install and configure the ODBC database drivers required to connect to the remote source. Each data source driver setup is described in its own section.
2. Create a remote source by selecting the appropriate adapter and configuring the connection properties and user credentials, using either SQL syntax or studio. If you use the generic ODBC adapter, familiarize yourself with configuration files and the generic adapter framework.  
See [Using the Generic Adapter Framework \[page 1239\]](#)
3. Create virtual tables to access the data stored in remote tables.  
See [Creating Virtual Tables \[page 1247\]](#)

## 9.1.1 Setting Up Database Drivers

The communication between SAP HANA and a remote data source is based on the ODBC protocol. To use the protocol, you need to install the appropriate drivers for the databases you want to connect to using SAP HANA smart data access.

### ODBC.INI File

SAP HANA smart data access requires an `.odbc.ini` file to be present in the administrator's home directory. Generally, an entry is created in the `.odbc.ini` file for each remote source. The `.odbc.ini` file may be empty, but it must be present.

### ODBC Driver Installation Location

ODBC driver library files need to be installed in a location that is searched by the SAP HANA server. If these libraries are placed in the SAP HANA `exe` directory, they will be found automatically. However, if they are installed elsewhere, you need to alter the `LD_LIBRARY_PATH` environment variable to point to this location. If this is not done, you may experience messages during SAP HANA smart data access queries stating the driver could not be loaded.

The `LD_LIBRARY_PATH` environment variable can be configured by creating or modifying the `.customer.sh` file in the home directory of the SAP HANA administrator user. This file should have a line that reads:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/path_to_driver_directory
```

- Since the `LD_LIBRARY_PATH` lists the directory, you do not need to specify a full directory path in the `.odbc.ini` `DSN, 'Driver='` entry or the `CREATE REMOTE SOURCE, 'Driver='` entry. These entries can simply specify the library name without the full path.
- You can validate that the changes in `.customer.sh` have taken effect by executing `echo $LD_LIBRARY_PATH` at the command prompt when logged in as the SAP HANA administrator.

### Example

If the IQ ODBC libraries are installed in `/opt/sybase/IQ-16_0/lib64`, the administrator's `$HOME/.customer.sh` file should contain a line that reads:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/opt/sybase/IQ-16_0/lib64
```

This statement adds the directory `/opt/sybase/IQ-16_0/lib64` to the end of the search path SAP HANA uses to find libraries.

## Remote SAP HANA Systems

When connecting to a remote SAP HANA system, SAP HANA Smart Data Access uses the SAP HANA ODBC driver installed with the SAP HANA server. Currently SAP HANA Smart Data Access does not support any HANA ODBC driver version other than the one installed by default with the SAP HANA server. Please check the SAP HANA driver requirements for the remote SAP HANA systems.

## Related Information

[SAP IQ Driver Setup \[page 1221\]](#)

[SAP ASE Driver Setup \[page 1221\]](#)

[SAP Event Stream Processor \(ESP\) Driver Setup \[page 1222\]](#)

[SAP Event Stream Processor \(ESP\) Connection Setup \[page 1223\]](#)

[SAP MaxDB Driver Setup \[page 1224\]](#)

[Teradata ODBC Version 13 Driver Setup \[page 1225\]](#)

[Teradata ODBC Version 14 Driver Setup \[page 1227\]](#)

[Microsoft ODBC Driver 11 for SQL Server 2012 Setup \[page 1228\]](#)

[Oracle Database 12C Setup \[page 1229\]](#)

[IBM DB2 Driver Setup \[page 1230\]](#)

[IBM Netezza Appliance Setup \[page 1231\]](#)

## 9.1.1.1 SAP IQ Driver Setup

Configure the SAP IQ ODBC driver to connect to an SAP IQ remote source.

### Prerequisites

You have downloaded and installed the SAP IQ ODBC driver from the SAP Software Download Center.

### Procedure

- DSN
  - a. Define a DSN entry in `.odbc.ini` for each SAP IQ remote source. For example:

```
[TESTIQ]
Driver= libdbodbc16_r.so
ServerName=testiq
CommLinks=tcip(host= server.com;port=2638)
```

- b. Define a remote source using a command such as:

```
create remote source TESTIQ adapter iqodbc configuration 'DSN=TESTIQ' with
CREDENTIAL TYPE 'PASSWORD' USING 'user=user;password=password';
```

- Driver properties
  - a. Specify all driver properties in the CREATE REMOTE SOURCE command (a DSN entry in `.odbc.ini` is not necessary). For example:

```
create remote source TESTIQ adapter iqodbc configuration
'Driver=libdbodbc16_r.so;ServerName=testiq;CommLinks=tcip(host=server.com;
port=2638)'with CREDENTIAL TYPE 'PASSWORD' USING
'user=user;password=password';
```

## 9.1.1.2 SAP ASE Driver Setup

Configure the SAP ASE ODBC driver to connect to an SAP ASE remote source.

### Prerequisites

You have downloaded and installed the SAP ASE ODBC driver from the SAP Software Download Center.

## Procedure

- DSN
  - a. Define a DSN entry in `.odbc.ini` for each SAP ASE remote source. For example:

```
[TESTASE]
Server=server.com
Port=4100
Driver= libsybdrvodb-sql11en8.so
Database=testdb
```

- b. Define a remote source using a command such as:

```
create remote source TESTASE adapter aseodbc configuration
'DSN=TESTASE'with CREDENTIAL TYPE 'PASSWORD' USING
'user=user;password=password';
```

- Driver properties
  - a. Specify all driver properties in the CREATE REMOTE SOURCE command (a DSN entry in `.odbc.ini` is not necessary). For example:

```
create remote source TESTASE adapter aseodbc configuration
'Server=server.com;Port=4100;Driver=libsybdrvodb-
sql11en8.so;Database=testdb' with CREDENTIAL TYPE 'PASSWORD' USING
'user=user;password=password';
```

### 9.1.1.3 SAP Event Stream Processor (ESP) Driver Setup

Install the ODBC driver for SAP Event Stream Processor (ESP) 5.1 SP04 or higher.

#### Prerequisites

You have installed the unixODBC driver, version 2.3.1 or higher, on the SAP HANA machine.

#### Context

Using SAP HANA Smart Data Access, you can access the content of an ESP window using simple SQL SELECT queries. A Linux ODBC driver lets you access the contents of ESP windows through the ODBC interface.

## Procedure

1. Launch the ESP installer.

2. When prompted for an installation location, select a writeable location on the machine hosting SAP HANA.
3. When prompted for an installation type, select *Custom*.
4. From the list of options, select only the *ODBC Driver for SAP Event Stream Processor* option.
5. When you have finished installing the driver, stop the SAP HANA instance you will be using with ESP.
6. On the machine hosting SAP HANA, open the `hdbenv.sh` file in the `DIR_INSTANCE` location.
7. Append the path to the ODBC driver to `LD_LIBRARY_PATH`.
8. Restart the SAP HANA instance.

## Related Information

[SAP Event Stream Processor \(ESP\) Connection Setup \[page 1223\]](#)

### 9.1.1.4 SAP Event Stream Processor (ESP) Connection Setup

Create a connection from an SAP HANA instance to SAP Event Stream Processor (ESP) 5.1 SP04 or higher.

#### Prerequisites

- `unixODBC`: The version of `unixODBC` needs to be 2.3.1 or higher on the machine where the SAP HANA instance is running.
- `LD_LIBRARY_PATH`: The `LD_LIBRARY_PATH` environment variable should include the following two directories from the ESP installation directory: `ESP_INSTALL/odbc` and `ESP_INSTALL/lib`.

#### Procedure

1. Create a remote source over the ESP window.

```
CREATE REMOTE SOURCE <source_name> ADAPTER "odbc" CONFIGURATION FILE
'property_esp.ini'
CONFIGURATION <connection_string> WITH CREDENTIAL TYPE 'PASSWORD' USING
'<credential_string>'
```

The format of the `<connection_string>` is as follows:

```
'Driver=<driver_path>;Database=<database_path>;Servername=<server_name>;Port=<port_number>;
Username=<username>;Password=<password>;SSLmode=disable;'
```

- Driver path: Path to the ESP driver `libesp_psqlodbc_a_lib.so`
- Database path: `<ESP project name>/<ESP workspace name>/user`

- Server name: Machine running the ESP instance
  - Port number: Port number of the ESP instance
  - User name/password: Credentials to access the ESP instance
2. Create a virtual table over the above remote source.

```
CREATE VIRTUAL TABLE <virtual_table_name> AT
<source_name>."<workspace_name>". "<project_name>". "<window_name>"
```

Replace `<workspace_name>`, `<project_name>`, and `<window_name>` with the names of the workspace, project, and the relevant window within the ESP instance that SAP HANA is connecting to. The table created above can be used like any other SAP HANA SDA virtual table.

### **i** Note

SAP HANA SDA 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 would either fail or produce incorrect data.

### Example

1. Create a remote source:

```
CREATE REMOTE SOURCE esp1 ADAPTER "odbc" CONFIGURATION FILE
'property_esp.ini' CONFIGURATION 'Driver=/sapmnt/HOME/i826198/'
```

2. Create a virtual table:

```
CREATE VIRTUAL TABLE etab1 AT esp1."ws1"."proj2"."WIN1"
```

3. Execute a query on the virtual table:

```
SELECT A."string1" FROM etab1 as A WHERE "int321" > 0
```

## Related Information

[SAP Event Stream Processor \(ESP\) Driver Setup \[page 1222\]](#)

### 9.1.1.5 SAP MaxDB Driver Setup

Install and configure the SAP MaxDB ODBC driver.

## Procedure

1. Download the SAP MaxDB package for Linux, `maxdb_all_linux_64bit_x86_64_7_9_08_x.tgz`, from the [SAP Store](#).

2. Extract the archive:

```
tar -xzvf maxdb_all_linux_64bit_x86_64_7_9_08_x.tgz
```

3. Unzip the MaxDB ODBC driver package SDBODBC.TGZ:

```
tar -xzvf maxdb-all-linux-64bit-x86_64-7_9_08_x/SDBODBC.TGZ
```

4. Create a MaxDB directory:

```
mkdir -p /opt/MaxDB/  
sudo chmod -R 777 /opt/MaxDB/
```

5. Copy the MaxDB ODBC driver directory to /opt/MaxDB:

```
cp -r lib /opt/MaxDB/
```

6. Add the DSN entry to the `odbc.ini` configuration file:

- a. Change to the `rxradm` home directory (`~/`):

```
vi ~/.odbc.ini
```

- b. Add the new entry as below:

```
[MaxDB]  
Driver=/opt/MaxDB/lib/libbdbodbcw.so  
Description=MaxDB  
ServerDB=YourServer  
SQLMode=INTERNAL  
Isolation Level =1
```

7. Verify the newly added DSN:

```
XIYL50837198A:/usr/sap/D70/HDB70> isql MaxDB USERNAME USERPASSWORD  
Connected!
```

8. `SQL> quit`

## Related Information

[Setting up the SAP MaxDB ODBC Driver](#)

### 9.1.1.6 Teradata ODBC Version 13 Driver Setup

Install and configure the Teradata ODBC Version 13 Driver.

## Prerequisites

Teradata drivers are packaged and distributed as a Linux RPM file. You can obtain these RPMs from Teradata.

## Procedure

1. Download and extract the following archives:

- `tdicu-13.10.00.00-1.tar.gz`
- `TeraGSS_suselinux-x8664__linux_x8664.13.10.00.06-1.tar.gz`
- `TeraGSS_redhatlinux-i386__linux_i386.13.10.00.06-1.tar.gz`
- `tdodbc__linux_x64.13.10.00.04-1.tar.gz`

2. Install the following extracted packages:

- `sudo rpm -i tdicu-13.10.00.00-1.noarch.rpm`
- `sudo rpm -i TeraGSS_suselinux-x8664-13.10.00.06-1.x86_64.rpm`
- `sudo rpm -i TeraGSS_redhatlinux-i386-13.10.00.06-1.i386.rpm`  
The installation of this package will normally fail but without it, the last RPM package won't install correctly because of missing dependencies.
- `sudo rpm -i tdodbc-13.10.00.04-1.noarch.rpm`

The driver will be installed in `/opt/Teradata`.

3. Change the default Kerberos 5 setup.

This is necessary because 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.

Make the required configuration changes as follows:

- a. Edit the `/opt/teradata/teragss/site/TdgssUserConfigFile.xml` file and add:

```
<Mechanism Name="KRB5">
  <MechanismProperties MechanismEnabled="no" />
</Mechanism>
```

- b. Remove the `/opt/teradata/teragss/site/linux-x8664/<version>/TdgssUserConfigFile.xml` file, if it exists.

- c. Run the following command:

```
/opt/teradata/teragss/linux-x8664/client/bin/run_tdgssconfig
```

4. To test the database use: `/opt/teradata/client/13.10/odbc_64/bin/tdxodbc`. You should consult the Teradata ODBC documentation, however, the information below describes a typical installation.

The `.odbc.ini` file **MUST** contain the `[ODBC]` and `[ODBC Data Source]` sections in addition to the DSN entries for each server:

```
[ODBC]
InstallDir=/opt/teradata/client/ODBC_64
[ODBC Data Sources]
default=tdata.so
TD=tdata.so
[TD]
Driver=/opt/teradata/client/ODBC_64/lib/tdata.so
DBCName=server.com
Username=
Password=
CharacterSet=UTF8
```

The matching CREATE REMOTE SOURCE statement for this is:

```
create remote source TD adapter tdodbc configuration 'DSN=TD' with CREDENTIAL
TYPE 'PASSWORD' USING 'user=user;password=password'
```

## 9.1.1.7 Teradata ODBC Version 14 Driver Setup

Install the Teradata ODBC Version 14 Driver.

### Prerequisites

Teradata drivers are packaged and distributed as a Linux RPM file. You can obtain these RPMs from Teradata.

### Procedure

1. Extract and install the following 3 packages:

- Linux\i386-x8664\tdicu\tdicu-14.00.00.00-1.noarch.rpm
- Linux\i386-x8664\TeraGSS\TeraGSS\_linux\_x64-14.00.00.00-1.noarch.rpm
- Linux\i386-x8664\tdodbc\tdodbc-14.00.00.02-1.noarch.rpm

The driver will be installed in `/opt/Teradata`.

2. Change the default Kerberos 5 setup.

This is necessary because 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.

Make the required configuration changes as follows:

- a. Edit the `/opt/teradata/teragss/site/TdgssUserConfigFile.xml` file and add:

```
<Mechanism Name="KRB5">
  <MechanismProperties MechanismEnabled="no" />
</Mechanism>
```

- b. Remove the `/opt/teradata/teragss/site/linux-x8664/<version>/TdgssUserConfigFile.xml` file, if it exists.
- c. Run the following command:

```
/opt/teradata/teragss/linux-x8664/client/bin/run_tdgssconfig
```

3. To test the database use: `/opt/teradata/client/14.00/odbc_64/bin/tdxodbc`. For more information, consult the Teradata ODBC documentation.

## 9.1.1.8 Microsoft ODBC Driver 11 for SQL Server 2012 Setup

Install and configure the Microsoft ODBC Driver 11 for Microsoft SQL Server 2012.

### Prerequisites

Before you install the Microsoft ODBC Driver 11 for SQL Server 2012, ensure that the workstation is clean.

### Procedure

1. Download the Microsoft ODBC Driver 11 for SQL Server 2012 from the Microsoft website.
2. Install the unixODBC 2.3.0 Driver Manager. Note that only version 2.3.0 is supported.
3. Install the Microsoft ODBC Driver 11 for SQL Server 2012.  
The driver installation creates the following directory: `/opt/microsoft/msodbcsql/lib64/libmsodbcsql-11.0.so.2260.0`.
4. Copy the `libmsodbcsql-11.0.so.2260.0` file (for SUSE Linux) or `libmsodbcsql-11.0.so.2270.0` file (for Red Hat Linux), to the admin directory, `/usr/sap/<sid>/HDB**/exe`.  
Note that your unixODBC libraries may be installed on `/usr/local/lib/` when you do not use an RPM package.
5. Add this directory to the `LD_LIBRARY_PATH`.  
Depending on your setup, you might also need to add `/usr/local/lib64` to the SAP HANA administrator user's `PATH` variable.
6. Open or create the `odbc.ini` file in your `/home` directory.
  - a. Add the following DSN entry:

```
[MSSQL]
Driver= /opt/microsoft/msodbcsql/lib64/libmsodbcsql-11.0.so.2260.0
Server=<host>,<port>
Database=<database_name>
```

    - Host: Machine running Microsoft SQL Server.
    - Port: Port number of Microsoft SQL Server (default 1433).
    - Database name: Name of the database created on Microsoft SQL Server (default master).
  - b. Copy the file to `/usr/sap/<sid>/HDB**/` and rename it `.odbc.ini`.

## 9.1.1.9 Oracle Database 12C Setup

Install and configure the Oracle Database Driver for Oracle Database 12c.

### Prerequisites

Before you install the Oracle Database driver, ensure that the workstation is clean.

#### **i** Note

When you use Oracle Database 12C, use the driver unixODBC version **2.3.2**.

### Procedure

1. Update unixODBC to version unixODBC-2.3.2.

After downloading the package, execute the following commands:

```
>./configure
>sudo make
>sudo make install
```

Use the `isql --help` command to check the version of unixODBC.

2. Install the Oracle Database 12c driver.
  - a. Download two packages: `instantclient-basic-linux.x64-12.1.0.1.0.zip` and `instantclient-odbc-linux.x64-12.1.0.1.0.zip`.
  - b. Directly unzip the driver packages after you download them.  
The folder `instantclient_12_1` will be generated. It will not conflict with other database libraries.
3. Create the Oracle TNS name and DSN entry (`.odbc.ini`):

Option	Description
Add the Oracle TNS entry to the <code>tnsnames.ora</code> file in <code>ADM***\$HOME</code>	<pre>ORCL= (DESCRIPTION = (AADDRESS = (PROTOCOL = TCP) (HOST = Hostname) (PORT = 1521) ) (CONNECT_DATA = (SERVER = DEDICATED) (SERVICE_NAME = SERVERNAME) ) )</pre>
Add the DSN section (for example, <code>ora12c</code> ) to the <code>.odbc.ini</code> file in <code>ADM*** \$HOME</code>	<pre>[ora12c] Driver=your_install_oracle_driver_folder/ instantclient_12_1/libsqora.so.12.1 ServerName=ORCL</pre>

4. Add the following environment variables to the `.customer.sh` file in the home directory of the SAP HANA administrator user:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:your_oracle_driver_dir/  
instantclient_12_1  
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib  
export TNS_ADMIN=~/
```

5. Restart SAP HANA.
6. Create the Oracle remote source with the following statement: `CREATE REMOTE SOURCE ORA_MY ADAPTER "odbc" CONFIGURATION FILE 'property_orcl.ini' CONFIGURATION 'DSN=ora12c' WITH CREDENTIAL TYPE 'PASSWORD' USING 'user=USER;password=PASSWORD;`

#### **i** Note

The Oracle remote sources do not support empty strings. Therefore, when values are inserted in a virtual table generated from an Oracle remote source, the empty string values are transformed into NULL values. This behavior also has an impact on some of the SDA optimization techniques (for example, join relocation).

## 9.1.1.10 IBM DB2 Driver Setup

SAP HANA supports IBM DB2 V10.1.

### Prerequisites

- You have installed the unixODBC 2.3.1 driver on the SAP HANA machine.
- You have downloaded the IBM DB2 driver from the IBM portal.

### Procedure

1. Install the IBM DB2 CLI/ODBC driver.
2. Add the DSN in the ODBC configuration file.
3. Verify the newly added DSN.

For more information about connecting to an IBM DB2 database system, see SAP Note [1382952](#).

## 9.1.1.11 IBM Netezza Appliance Setup

SAP HANA supports IBM Netezza 7.

### Prerequisites

- You have installed the unixODBC 2.3.0 driver on the SAP HANA machine. Note that only version 2.3.0 is supported.
- You have downloaded the IBM Netezza driver from the IBM portal.

### Procedure

1. Install the IBM Netezza driver.
2. Configure the `.odbc.ini` file.
3. Configure the `.odbcinst.ini` file.
4. Verify your setup.

### Next Steps

Due to a Netezza-specific limitation, real cursors are not supported. (A driver configuration that can be set to enable cursors does not exist.) The workaround is to set a high value for the Prefetch Count in the `.odbcinst.ini` file as follows:

```
[NetezzaSQL]
Driver = /usr/nz70/lib64/libnzodbc.so
Setup = /usr/nz70/lib64/libnzodbc.so
APILevel = 1
ConnectFunctions = YYN
Description = Netezza ODBC driver
DriverODBCVer = 03.51
DebugLogging = true
LogPath = /tmp
UnicodeTranslationOption = utf16
CharacterTranslationOption = all
PreFetch = 25600
Socket = 16384
```

## 9.1.2 Creating and Configuring Remote Data Sources

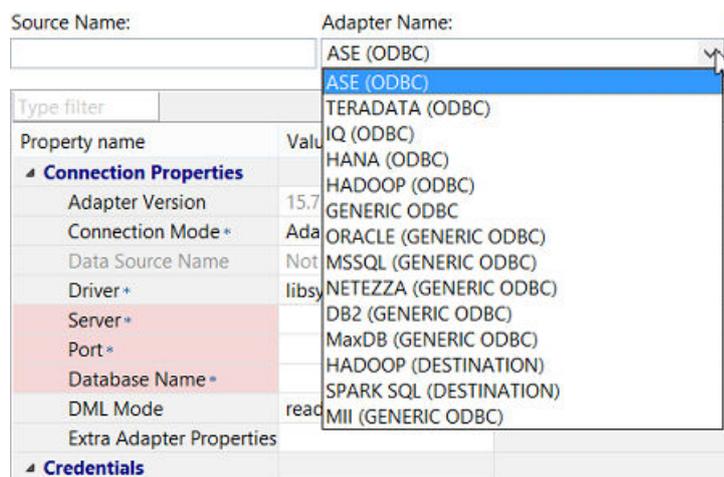
You can create and configure remote sources using the SAP HANA studio remote source editor or the CREATE REMOTE SOURCE and ALTER REMOTE SOURCE commands.

### Prerequisites

You require the system privilege CREATE REMOTE SOURCE or DATA ADMIN to set up remote sources.

### Adapter Types

There are dedicated and generic ODBC adapters available for smart data access, as shown below:



#### Dedicated Adapters

SAP HANA has built-in native specialized code to support specific data sources, such as SAP ASE, SAP IQ, and Teradata Database.

#### Generic Adapters

Instead of built-in specialized code, a configuration file is used to specify the parameters specific to the database being used:

- A template configuration file is provided in the `DIR_EXECUTABLE` location of the SAP HANA installation for specific data sources supported through the generic ODBC adapter, such as Microsoft SQL Server and Oracle Database.
- The generic ODBC adapter also allows you to connect to databases not included in the list of supported remote sources. Since there is no template configuration file provided for these cases, you have to manually create a configuration file from scratch based on the ODBC driver that you are using.

For more information, see *Using the Generic Adapter Framework*.

## Connection Modes

There are two ways in which you can connect to the remote data source:

- Using the adapter properties: You directly specify the driver installed on the SAP HANA instance. You need to give the full connection string.
- Using the data source name (DSN): You specify the DSN defined on the SAP HANA instance.

Which options are available depends on the remote source being used. If you are using the remote source editor, explicitly select the connection mode to enter the connection information required by the chosen ODBC adapter:

Remote Source Editor	SQL Statement																						
<p>Source Name: <input type="text"/> Adapter Name: ASE (ODBC) ▼</p> <p>Type filter <input type="text"/></p> <table border="1"><thead><tr><th>Property name</th><th>Value</th></tr></thead><tbody><tr><td colspan="2">Connection Properties</td></tr><tr><td>Adapter Version</td><td>15.7</td></tr><tr><td>Connection Mode*</td><td>Adapter Properties ▼</td></tr><tr><td>Data Source Name</td><td>Data source name</td></tr><tr><td>Driver*</td><td>Adapter Properties</td></tr><tr><td>Server*</td><td></td></tr><tr><td>Port*</td><td></td></tr><tr><td>Database Name*</td><td></td></tr><tr><td>DML Mode</td><td>readonly</td></tr><tr><td>Extra Adapter Properties</td><td></td></tr></tbody></table>	Property name	Value	Connection Properties		Adapter Version	15.7	Connection Mode*	Adapter Properties ▼	Data Source Name	Data source name	Driver*	Adapter Properties	Server*		Port*		Database Name*		DML Mode	readonly	Extra Adapter Properties		<p>Adapter properties:</p> <p>Sample Code</p> <pre>CREATE REMOTE SOURCE TESTASE ADAPTER "aseodbc" CONFIGURATION 'Server=server.com;Port=4100; Driver=libsybdrvodb-sql1en8.so; Database=testdb' WITH ...</pre> <p>DSN:</p> <p>Sample Code</p> <pre>CREATE REMOTE SOURCE TESTASE ADAPTER "aseodbc" CONFIGURATION 'DSN=TESTASE' WITH ...</pre>
Property name	Value																						
Connection Properties																							
Adapter Version	15.7																						
Connection Mode*	Adapter Properties ▼																						
Data Source Name	Data source name																						
Driver*	Adapter Properties																						
Server*																							
Port*																							
Database Name*																							
DML Mode	readonly																						
Extra Adapter Properties																							

## Related Information

[Create Remote Data Sources \[page 1234\]](#)

[Using the Generic Adapter Framework \[page 1239\]](#)

## 9.1.2.1 Create Remote Data Sources

A remote data source allows you to connect to another source or system and remotely access its data without having to replicate this data into SAP HANA.

### Prerequisites

- You have configured an ODBC connection from SAP HANA to the remote database.
- The remote data source is reachable by the network from the computer you are using.

### Context

The procedure described below uses the SAP HANA studio. As an alternative, you can create a remote source through SQL using the `CREATE REMOTE SOURCE` command.

### Procedure

1. In the SAP HANA studio in the *Systems* view, expand the *Provisioning* node.
2. Select the *Remote Sources* folder and from the context menu choose *New Remote Source*.

The *New Remote Source* tab appears.

Source Name:	Adapter Name:	Source Location:
	ASE (ODBC)	indexserver
Type filter		
Property name	Value	
▲ Connection Properties		
Adapter Version	15.7	
Connection Mode *	Adapter Properties	
Data Source Name	Not applicable	
Driver *	libsybdrvodb-sqlen8.so	
Server *		
Port *		
Database Name *		
DML Mode	readonly	
Extra Adapter Properties		
▲ Credentials		
Credentials Mode *	Technical user	
User Name *		
Password *		

3. In the *Adapter Name* dropdown list, choose the appropriate adapter (MSSQL is chosen in the example below).
4. Enter the following information:
  - Source name: The name of your remote source

- Connection details: The connection properties as required by the selected adapter type, including the DML mode (**readwrite** or **readonly** (default value)).
- User credentials:
  - Technical user: All connections to the remote data source share the same credential for the data source.
  - Secondary credentials: There is one credential per user per data source.
  - SSO (Kerberos): All connections to the remote source (SAP HANA remote sources only) are authenticated through Kerberos single sign-on (SSO).

For more information, see *SAP HANA Smart Data Access* in the *SAP HANA Security Guide*.

Example:

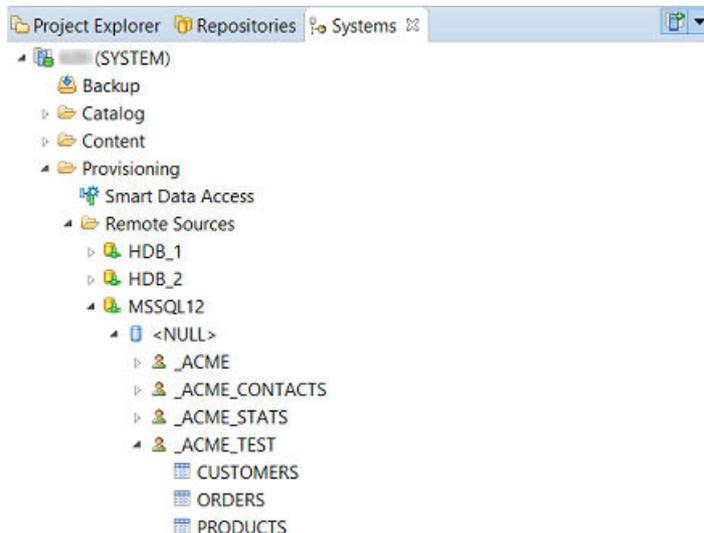
Source Name:	Adapter Name:	Source Location:
MSSQL12	MSSQL (GENERIC ODBC)	indexserver

Property name	Value
<b>Connection Properties</b>	
Adapter Version	SQL Server 2012
Connection Mode	Data source name
Configuration file	property_mss.ini
Data Source Name	mssql
DML Mode	readonly
<b>Credentials</b>	
Credentials Mode	Technical user
User Name	*****
Password	*****

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.



## Related Information

[Enable Read-Write Access to a Remote Source \[page 1236\]](#)

## 9.1.2.2 Enable 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.

### 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.

### Procedure

- Set the DML mode in the SAP HANA Studio remote source editor:

In the *Systems* view, expand the **Provisioning > Remote Sources** node. Select the remote source and from the context menu choose *Open Definition*. You can now view and change the value of the *DML Mode* field.

- Set the DML mode using a DDL command as shown below:

```
CREATE REMOTE SOURCE TEST_HANA_LOOPBACK ADAPTER "hanaodbc" CONFIGURATION
'ServerNode=hanaserver:30015;Driver=libodbcHDB.so;
DML_MODE=READONLY' WITH CREDENTIAL TYPE 'PASSWORD' USING
'user=SYSTEM;password=manager'
```

- Find out which DML mode option has been set using the following query:

```
SELECT SUBSTR_AFTER (CONNECTION_INFO, 'dml_mode">' ) "DML Mode" from
PUBLIC.REMOTE_SOURCES
where REMOTE_SOURCE_NAME= 'UK_ORDERS';
```

## 9.1.2.3 Alter a Remote Source to Use the Data Provisioning Server

Transform a Smart Data Access remote source into a Smart Data Integration data source.

### Prerequisites

- The Data Provisioning Server (DP Server) has been started. For more information, see [Enable the Data Provisioning Server](#).
- The DP adapters have been configured. For more information, see [Data Provisioning Adapters](#).

#### **i** Note

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 \[page 1360\]](#).

### Context

You can use the `ALTER REMOTE SOURCE` command to alter a remote source so that it uses the DP Server. The following conversions are supported:

SDA Adapter	DP Server Adapter
MSSQL (GENERIC ODBC)	MssqlLogReaderAdapter
ORACLE (GENERIC ODBC)	OracleLogReaderAdapter
DB2 (GENERIC ODBC)	DB2LogReaderAdapter
TERADATA (ODBC)	TeradataAdapter

Note that you can also use the `ALTER REMOTE SOURCE` feature to perform the conversion in the inverse direction, that is, to convert a Smart Data Integration (DP) remote source into a Smart Data Access remote source.

The procedure below shows how to edit the remote source in the SAP HANA studio.

### Procedure

1. Open the existing remote source configuration. For example:

Source Name:	Adapter Name:	Source Location:
MSSQL12	MSSQL (GENERIC ODBC)	indexserver

Property name	Value
<b>Connection Properties</b>	
Adapter Version	SQL Server 2012
Connection Mode +	Data source name
Configuration file +	property_mss.ini
Data Source Name +	mssql
DML Mode	readonly
<b>Credentials</b>	
Credentials Mode +	Technical user
User Name +	*****
Password +	*****

2. Select the new adapter type and enter the new configuration details. For example:

Source Name:	Adapter Name:	Source Location:
MSSQL12	MssqlLogReaderAdapter	agent (DemoAgent)

Property name	Value
<b>Configurations</b>	
<b>Generic</b>	
Instance Name +	
Administration Port +	
<b>Database</b>	
Data Server +	
Port Number +	
DAC Port Number +	
Database Name +	
Database Data Capture Mode +	Native Mode
Use Remote Database +	false
<b>Security</b>	
Use SSL	false
Host Name in Certificate	Not applicable
<b>LogReader</b>	
Ignore log record processing errors	false
<b>Credentials</b>	
Credentials Mode +	Technical user
<b>Credential</b>	
User Name	Not applicable
Password	Not applicable

3. Save your changes.

## Results

The remote source has been converted into an SDI (DP) data source. All existing virtual tables of the remote source have also been altered.

## 9.1.2.4 Delete Remote Sources

You have an existing remote source in the SAP HANA *Systems* view that you want to remove.

### Procedure

1. In the *Systems* view, expand the **Provisioning > Remote Sources** node.
2. Select the remote source and from the context menu choose *Delete*.
3. Choose *Yes* to confirm.

### Results

The remote source disappears from the *Systems* view, including all dependent virtual tables.

## 9.1.3 Using the Generic Adapter Framework

To enable SAP HANA to interoperate with other database-like sources that support the ODBC interface, the generic adapter framework provides a configurable way of specifying the behavior and capabilities of ODBC data sources in a way that does not require specialized native support to be built into SAP HANA.

This is necessary because SAP HANA has built-in native specialized code to support only a few select data source types, for example, SAP ASE, SAP IQ, Teradata Database, and so on. The main reason why a wider set of ODBC-based data sources is not supported is that some of them have configuration and behavior that cannot be handled purely through the ODBC interface.

The generic adapter framework therefore allows you to specify an explicit properties text file that defines the capabilities, function signatures, data type mappings and additional properties specific to each instance of a remote data source that is created in SAP HANA.

### When to Use the Generic Adapter Framework

Find out whether you need to use the generic adapter framework as follows:

1. Check whether there is a native adapter available for the data source in question, for example, SAP ASE or SAP IQ.
  - If there is, you can directly create the remote source in SAP HANA using the correct adapter label with the CREATE REMOTE SOURCE command, for example, as follows:

```
CREATE REMOTE SOURCE ASE1 ADAPTER "ASEODBC" ...
```

- If there isn't, go to step 2.
- 2. Check whether a template properties file is provided with the SAP HANA installation for the data source in question, for example, Microsoft SQL Server or Oracle Database.  
A sample properties file for the Oracle data source (`property_orcl.ini`) and Microsoft SQL Server data source (`property_mss.ini`) is provided in the `DIR_EXECUTABLE` location of each SAP HANA installation.
  - If there is a template, create a properties file, for example, `userx_oracle1.ini`, based on this template. If there are features that you want to disable or override, edit the properties file accordingly. Make sure it is correct.  
Create the remote source in SAP HANA using the generic adapter and the properties file created above.
  - If there isn't, go to step 3.
- 3. As there is no template properties file provided for the type of data source in question, you have to manually create a properties file from scratch based on the ODBC driver that you are using and the capabilities, functions, type names, and so on that the driver supports.  
Note that you only need to add values that are different from the default values.

### Note

Incorrect values specified in the properties file could cause incorrect results and incorrect behavior in the server.

A properties file when specified is only for the specific instance of the data source created and not for a generic category of data source types. The same properties file could be specified for multiple data source instances of the same type.

### Example

Create a remote source using the properties file:

```
CREATE REMOTE SOURCE Ora1 adapter "ODBC" CONFIGURATION FILE 'oraprops1.ini'  
CONFIGURATION 'Server=testoracle;Port=4100...'
```

```
CREATE REMOTE SOURCE Ora2 adapter "ODBC" CONFIGURATION FILE 'oraprops1.ini'  
CONFIGURATION 'DSN=production_oracle;...'
```

## Related Information

[Properties File \[page 1241\]](#)

[Sample Properties File \[page 1242\]](#)

## 9.1.3.1 Properties File

A properties text file defines the capabilities, function signatures, data type mappings, and additional properties specific to each instance of a remote data source that is created in SAP HANA.

### Properties File Syntax

Each line entry in the properties file has the following format:

```
propname : value
```

Name and value are string literals, where name has the following form:

```
propname : CAP_name | FUNC_name | TYPE_name | PROP_name
```

- Entries beginning with "CAP\_" are interpreted as a CAPABILITY specification. These specify support for features or operators, such as JOINS, SUBQUERIES, and so on.
- Entries beginning with "FUNC\_" are interpreted as a function signature or support specification. These specify how functions in SAP HANA map to functions on the remote data source.
- Entries beginning with "TYPE\_" are interpreted as a type name mapping specification. These specify what the data type names are for the data types in SAP HANA.
- Entries beginning with "PROP\_" are interpreted as a property specification. These are other miscellaneous properties that are understood by SAP HANA.

#### Example

For a data source requiring the UNIX ODBC manager, the properties below could be defined as follows:

```
CAP_name : "true" | "false"  
FUNC_name : "true" | "false" | func_string_expression  
TYPE_name : type_name_expression  
PROP_name : prop_value_expression
```

- `func_string_expression` can include the special words \$1, \$2, \$3 ..., which will be expanded to the first, second, third (and so on) arguments of the function. \$\* will be expanded as a comma-separated list of all arguments passed to the function.
- `type_name_expression` can include the special words \$PRECISION and \$SCALE, which will be expanded to the precision and scale of the type name.

## CREATE REMOTE SOURCE Syntax

Create the remote source using the following syntax:

```
CREATE REMOTE SOURCE srcname ADAPTER "ODBC" CONFIGURATION FILE 'filename'  
CONFIGURATION str_const [opt_credentials_clause]
```

In this example, `filename` is the name of the file containing the properties specific to the remote data source being created. SAP HANA will look for this file name in the `DIR_EXECUTABLE` environment variable location, which is currently also used for other `.ini` files.

### **i** Note

The properties file can only be specified for the ADAPTER "ODBC", which specifies that the data source is a generic data source supporting the ODBC interface.

When the remote data source is created, the associated properties file is parsed and interpreted. All capabilities, function mappings, type name mappings, and property specifications are associated with the data source and used for all interoperations between SAP HANA and this source. For specifications that are not in the properties file, the generic ODBC interface is used:

- The set of definable properties is fixed based on the existing infrastructure in SAP HANA.
- It is explicitly specified to be extendible in the future based on future needs but would be backward compatible.

## Related Information

[Sample Properties File \[page 1242\]](#)

### 9.1.3.2 Sample Properties File

The following is a fragment of a properties file that could be used when creating a remote data source over an ASE instance.

```
CAP_SUBQUERY : true
CAP_SUBQUERY_GROUPBY : true
CAP_SUBQUERY_SELLIST : true
CAP_SUBQUERY_UPDATE : true
CAP_ORDERBY : true
CAP_ORDERBY_EXPRESSIONS : true
CAP_ORDERBY_UNRELATED : true
CAP_ORDERBY_UPDATE : true
CAP_JOINS : true
CAP_JOINS_OUTER : true
CAP_JOINS_MIXED : true
CAP_GROUPBY : true
CAP_GROUPBY_ALL : true
CAP_AGGREGATES : true
CAP_AGGREGATE_COLNAME : true
CAP_AND : true
CAP_OR : true
CAP_LIKE : true
CAP_LIKE_TSQL : true
FUNC_ABS : true
FUNC_ACOS : true
FUNC_ADD : true
FUNC_ADD_DAYS : DATEADD(DAY, $2, $1)
FUNC_ADD_MONTHS : DATEADD(MONTH, $2, $1)
FUNC_ADD_SECONDS : DATEADD(SECOND, $2, $1)
FUNC_ADD_YEARS : DATEADD(YEAR, $2, $1)
```

```
FUNC_ASCII : true
FUNC_ASIN : trueFUNC_ATAN : true
FUNC_TO_VARBINARY : false
FUNC_TO_VARCHAR : custom
FUNC_TRIM_BOTH : LTRIM(RTRIM($1))
FUNC_TRIM_LEADING : LTRIM($1)
FUNC_TRIM_TRAILING : RTRIM($1)
FUNC_UMINUS : false
FUNC_UPPER : true
FUNC_WEEKDAY : false
TYPE_TINYINT : tinyint
TYPE_ST_GEOMETRY : image
TYPE_LONGBINARY : image
TYPE_LONGCHAR : text
TYPE_DATE : datetime
TYPE_TIME : datetime
TYPE_DATETIME : datetime
TYPE_REAL : real
TYPE_SMALLINT : smallint
TYPE_INT : int
TYPE_FLOAT : float
TYPE_CHAR : text
PROP_USE_UNIX_DRIVER_MANAGER : true
```

## 9.1.4 Setting Up Single Sign-On (SSO) with Kerberos

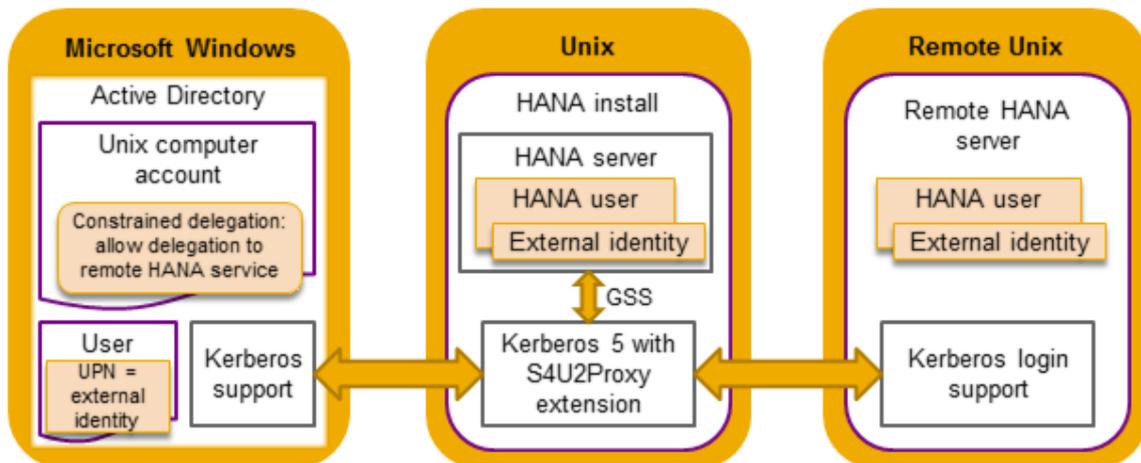
SAP HANA smart data access supports single sign-on with Kerberos for connections to SAP HANA remote sources. Using Kerberos constrained delegation and protocol transition, it allows SAP HANA users to be authenticated automatically on Microsoft Windows Active Directory, without having to provide a password (SSO mode).

### Prerequisites

Microsoft Windows Server, at least version 2003.

## 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>&lt;sidadm home&gt;/etc/krb5_hdb.conf</code>	Configuration of the Kerberos realm to be used with the SAP HANA server installed under <code>&lt;sidadm&gt;</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>
```

## Tasks

On the source SAP HANA server, that is, where SDA is used, you need to configure Kerberos to support constrained delegation. The high-level steps are as follows:

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` 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.
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.

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 (*SAP HANA Smart Data Access: SSO with Kerberos and Microsoft Windows Active Directory*).

## Related Information

[SAP Note 2303807](#) 

[Creating Remote Sources with Kerberos Authentication \[page 1246\]](#)

## 9.1.4.1 Creating Remote Sources with Kerberos Authentication

The Kerberos credential type allows users connecting to SAP HANA remote sources to be authenticated through single sign-on (SSO).

You can configure a remote source with Kerberos authentication in the following ways:

- From the **New Remote Source UI**, for example, in the SAP HANA studio  
Open the *New Remote Source* UI (expand the *Provisioning* node, select the *Remote Sources* folder, and from the context menu choose *New Remote Source*).  
In the *Credentials Mode* field in the *Credentials* section, select *SSO (Kerberos)*, as shown in the example below. You do not specify a user or password:

Source Name:	Adapter Name:
loopback	HANA (ODBC)
Type filter	
Property name	Value
▼ Connection Properties	
Adapter Version	1.0
Connection Mode *	Adapter Properties
Data Source Name	Not applicable
Driver *	libodbcHDB.so
Server *	localhost
Port *	30015
DML Mode	readonly
Extra Adapter Properties	
▼ Credentials	
Credentials Mode *	SSO (Kerberos)
User Name	Not applicable
Password	Not applicable

- Using the **CREATE REMOTE SOURCE** SQL statement  
Use the CREATE REMOTE SOURCE statement together with the parameter WITH CREDENTIAL TYPE 'KERBEROS'. You do not specify a user or password. The syntax is shown below:

```
CREATE REMOTE SOURCE <remote_source_name> ADAPTER <adapter_name>  
CONFIGURATION <connection_info_string> WITH CREDENTIAL TYPE 'KERBEROS'
```

### Sample Code

```
CREATE REMOTE SOURCE "H1" ADAPTER "hanaodbc" CONFIGURATION  
'Driver=libodbcHDB.so;ServerNode=hanaserver:3<ID>15;' WITH CREDENTIAL TYPE  
'KERBEROS';
```

- Using the **CREATE CREDENTIAL** SQL statement  
Use the CREATE CREDENTIAL statement together with the parameter TYPE 'KERBEROS'. The syntax is shown below:

```
CREATE | ALTER CREDENTIAL FOR [USER <user_name>] COMPONENT <component_id>  
PURPOSE <purpose_def> TYPE <type_def> [USING <using_param>]
```

### Sample Code

```
CREATE CREDENTIAL FOR COMPONENT 'SAPHANAFEDERATION' PURPOSE 'H1' TYPE  
'KERBEROS';
```

The KERBEROS credential type can be declared either as a global credential type for the remote source or as the individual type for a given user:

- All users can set KERBEROS as the authentication type if they want, since they are allowed to modify their own user credentials.
- User credentials are used only if there is no global credential defined for a PURPOSE (remote source in SDA scenarios).

Note that only one global credential type is allowed for a given PURPOSE (remote source). Similarly, there should be only one credential type per USER and PURPOSE. If multiple credential types are defined, the first one to be retrieved from the credential store will be used.

### **i** Note

The Kerberos credential type is intended to be used mainly as an alternative to the technical user credential type.

## Related Information

[Create Remote Data Sources \[page 1234\]](#)

## 9.1.5 Creating Virtual Tables

Virtual tables represent the tables in the remote data source.

You can create virtual tables in the SAP HANA studio as follows:

Option	Description
<b>By remote object (source)</b>	Select a remote object in the <i>Provisioning</i> section of the SAP HANA studio, then browse through the Catalog and select the schema where it should be added as a virtual table.  See: <a href="#">Create Virtual Tables by Remote Object [page 1248]</a>
<b>By schema (target)</b>	Select a target schema in the Catalog of the SAP HANA studio, then browse through <i>Provisioning</i> and select the remote object to be added as a virtual table.  See: <a href="#">Create Virtual Tables by Schema [page 1249]</a>

## 9.1.5.1 Create Virtual Tables by Remote Object

In the *Provisioning* section of the SAP HANA studio, select the remote object to be added as a virtual table, then browse through the Catalog and select the target schema.

### Prerequisites

You have already created a remote source. It appears in the *Systems* view under **Provisioning > Remote Sources**.

### Procedure

1. In the *Systems* view, expand the remote source to see the users and tables: **Provisioning > Remote Sources > <remote-source> > <user>**.
2. Select the remote object on which you want to create the virtual table and from the context menu choose *Add as Virtual Table*.  
The *Create Virtual Table* dialog box appears.
3. Enter a table name and select the target schema from the dropdown list.
4. Choose *Create*.  
An information box appears confirming that the virtual table has been created in the specified schema.

### Results

The new virtual table appears in the *Systems* view under **Catalog > <schema> > Tables**. The  icon indicates that it is a virtual table. Select the table and from the context menu choose *Open Definition*. You can see that the table type is *Virtual*.

Table Name:	Schema:	Type:	Source Name:	Remote Object:
MSSQL12_PRODUCT	SDA	Virtual	MSSQL12	PRODUCTS@<NULL>

Columns	Comment					
Name	SQL Data Type	Di...	Column Store Data...	Key	Not Null	Default
1 COLUMN_0	NVARCHAR	9		X(1)	X	
2 COLUMN_1	NVARCHAR	8				

## 9.1.5.2 Create Virtual Tables by Schema

Select the target schema in the Catalog of the SAP HANA studio, then browse through *Provisioning* and select the remote object to be added as a virtual table.

### Prerequisites

You have already created a remote source. It appears in the *Systems* view under **► Provisioning ► Remote Sources ►**.

### Procedure

1. In the *Systems* view, expand the folders **► <system> ► Catalog ► <schema> ►**.
2. Select the *Tables* folder and from the context menu choose *New Virtual Table*.

The *New Virtual Table* tab appears:

- The *Schema* field contains the name of the schema you selected.
- The *Type* field defines the table as a virtual table.

Table Name:	Schema:	Type:	Source Name:	Remote Object	
<input type="text"/>	SDA ▼	Virtual ▼	<input type="text"/>	<input type="text"/>	<input type="button" value="Browse..."/>

3. Select the remote object:
  - a. Choose *Browse*.  
A dialog box appears.
  - b. Select the relevant remote source and drill down to the relevant table.
  - c. Select the remote object and choose *OK*.

The source name and the remote object fields are filled in automatically with the names of the selected components. The virtual table name is constructed as follows and can be overwritten: `<remote-source>_<remote-object>`

4. Correct the target schema, if necessary.
5. Choose the *Save the Editor* icon in the upper right corner of the screen.

## Results

The new virtual table appears in the *Systems* view under **Catalog > <schema> > Tables**. The  icon indicates that it is a virtual table:

- ▾  SDA
  -  Column Views
  -  EPM Models
  -  EPM Query Sources
  -  Functions
  -  Indexes
  -  Procedures
  -  Sequences
  -  Synonyms
  - ▾  Tables
    -  MSSQL12\_PRODUCTS MSSQL12
  -  Triggers
  -  Views

### 9.1.5.3 Delete Virtual Tables

You have an existing virtual table in the SAP HANA *Systems* view that you want to remove.

#### Procedure

1. In the *Systems* view, expand the nodes **Catalog > <schema> > Tables**.
2. Select the virtual table and from the context menu choose *Delete*.

A warning appears if there are dependent virtual tables, stating that they will also be deleted. A list of these dependent tables is displayed.

3. Choose *OK* to confirm.

## Results

The virtual table disappears from the *Systems* view.

## 9.1.6 Creating Statistics on Virtual Tables

Statistics are one of the key facts that assist the query optimizer in making better decisions.

In particular, for some optimizations like join-relocation or semi-join rules, it is important to have accurate table or columns statistics in order to decide whether or not they should be applied, especially to avoid transferring too many rows from the remote source if it is not necessary.

You can create statistics with the following SQL command:

```
CREATE STATISTICS [<data_statistics_name>] ON <data_sources>
 [<data_statistics_properties>]
 ... <data_statistics_properties> := <data_statistics_property>...
 <data_statistics_property> ::= TYPE <data_statistics_type>
 <data_statistics_type> ::= HISTOGRAM | SIMPLE | ALL | RECORD COUNT
```

For more information about this command, see *CREATE STATISTICS* in the *SAP HANA SQL and System Views Reference*.

### Statistics Types

#### SIMPLE and HISTOGRAM Statistics

In order to evaluate the costs of semi-join optimizations, you should create simple statistics on all fields that are potentially in a join condition. SIMPLE statistics provide min, max, null count, count, and distinct count values. However, a full set of statistics also includes histogram information, which typically causes a higher workload than simple statistics.

#### RECORD COUNT Statistics

RECORD COUNT is available for virtual tables only and specifies that only the number of records is computed. This statistics type should take much less time to compute compared to the SIMPLE and HISTOGRAM types. However, in terms of query execution, the query optimizer has less information, which could lead to less optimized query execution. You should therefore only use the RECORD COUNT type if it is too expensive to create SIMPLE or HISTOGRAM statistics.

Note that RECORD COUNT statistics can be created on tables only, not on columns.

### Statistics Query Monitoring

You can check that the CREATE STATISTICS statement has been correctly executed using the Smart Data Access remote statements monitoring tool.

The example below shows that statistics queries (type SIMPLE) have been executed for each column of the virtual table:

The screenshot shows the 'Remote Statements Monitor' interface. A table lists several SQL statements, each with a timestamp and a status. A pop-up window titled 'Full SQL Statement' displays the following query:

```
SELECT MIN("MSSQL12_PRODUCTS"."COLUMN_20"),
MAX("MSSQL12_PRODUCTS"."COLUMN_20"), COUNT(*),
COUNT(DISTINCT "MSSQL12_PRODUCTS"."COLUMN_20"),
COUNT("MSSQL12_PRODUCTS"."COLUMN_20") FROM
"MYSCHEMA"."PRODUCTS" "MSSQL12_PRODUCTS"
```

## Statistics Results

Virtual table statistics are stored in the SYS"."P\_STATISTICS\_" system table. The example below shows the simple statistics values computed for each column of the specified table:

```
SELECT TOP 1000 * FROM "SYS"."P_STATISTICS_"
```

	SCHEMA...	TABLE_NAME	COLUMN_NA...	P...	MINVALUE	MINVAL...	MAXVALUE	MAXV...	COUNT	DCOUNT	NULLCOUNT
1	SDA	MSSQL12_PRODUCTS	COLUMN_0	1	18.370.83...	AD-1000	22.609.45...	PS-10...	108	108	0
2	SDA	MSSQL12_PRODUCTS	COLUMN_1	2	18.370.64...	AD	23.742.09...	TYPEC...	108	3	0
3	SDA	MSSQL12_PRODUCTS	COLUMN_2	3	18.688.81...	Beamer	24.610.86...	Workst*	108	26	0
4	SDA	MSSQL12_PRODUCTS	COLUMN_3	4	13.563.78...	0000000	20.346.82...	HISTO...	108	2	0
5	SDA	MSSQL12_PRODUCTS	COLUMN_4	5	14.126.73...	2012-10	20.346.82...	HISTO...	108	2	0
6	SDA	MSSQL12_PRODUCTS	COLUMN_5	6	13.563.78...	0000000	20.346.82...	HISTO...	108	2	0
7	SDA	MSSQL12_PRODUCTS	COLUMN_6	7	14.126.73...	2012-10	20.346.82...	HISTO...	108	2	0
8	SDA	MSSQL12_PRODUCTS	COLUMN_7	8	13.845.25...	1000000	22.026.84...	NAME...	108	107	0
9	SDA	MSSQL12_PRODUCTS	COLUMN_8	9	13.845.25...	1000000	19.216.52...	DESCID	108	107	0
10	SDA	MSSQL12_PRODUCTS	COLUMN_9	10	13.564.88...	0100000	23.456.22...	SUPPL...	108	46	0

In the case of the statistics type RECORD COUNT, only the column COUNT has a significant value, as shown in the example below:

```
SELECT TOP 1000 * FROM "SYS"."P_STATISTICS_"
```

	SCHEMA...	TABLE_NAME	COLUMN_NA...	P...	MINVALUE	MAXVALUE	COUNT	DCOUNT	NULLCOUNT
1	SDA	MSSQL12_PRODUCTS		0	-1	-1	108	-1	-1

Note that you can use the DROP STATISTICS statement to drop existing statistics:

```
DROP STATISTICS {<data_statistics_name> | ON <data_sources> [TYPE {HISTOGRAM | SIMPLE | ALL}]}
```

## Related Information

[Monitor Remote Connections and Remote Statements \[page 1253\]](#)

### 9.1.6.1 Statistics Retrieval from the Remote Source

For certain remote sources, the statistics for virtual tables can be retrieved by querying a remote table.

This functionality is supported only for the following remote sources: SAP HANA, SAP IQ, and Teradata Database. The name of the remote virtual table that is queried for statistics retrieval is "SYSTEM". "SDA\_STATISTICS" for SAP HANA, and "SYS\_STATISTICS" for SAP IQ and Teradata. "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
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)

## 9.1.7 Monitor Remote Connections and Remote Statements

Both the SAP HANA studio and SAP HANA cockpit provide monitoring functions for Smart Data Access.

### Prerequisites

You have the privileges granted by the role `sap.hana.admin.roles::Monitoring`.

## Context

You can 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.

## Procedure

Open the SAP HANA Smart Data Access Administration tools as follows:

Option	Procedure
SAP HANA studio	<ol style="list-style-type: none"><li>1. In the SAP HANA studio in the <i>Systems</i> view, expand your system's <i>Provisioning</i> node.</li><li>2. Select <i>Smart Data Access</i> and from the context menu choose <i>Open Smart Data Access Administration</i>.</li><li>3. Choose the appropriate tab:<ul style="list-style-type: none"><li>◦ Query Monitoring</li><li>◦ Connection Monitoring</li></ul></li></ol>
SAP HANA cockpit	<ol style="list-style-type: none"><li>1. In the SAP HANA studio in the <i>Systems</i> view, select your system and from the context menu choose ► <i>Configuration and Monitoring</i> ► <i>Open SAP HANA Cockpit</i> 🗒️. Note that you can also open the SAP HANA cockpit with the following URLs:<ul style="list-style-type: none"><li>◦ <code>https://&lt;host_FQDN&gt;:43&lt;instance&gt;/sap/hana/admin/cockpit</code> (recommended)</li><li>◦ <code>http://&lt;host_FQDN&gt;:80&lt;instance&gt;/sap/hana/admin/cockpit</code></li></ul>The <i>Remote Connections Monitor</i> and <i>Remote Statements Monitor</i> tiles are displayed under <i>Smart Data Access Administration</i>. Note that if the tiles are not visible, you can add them from the tile catalog.</li><li>2. Click the appropriate tile to launch the monitoring app.</li></ol>

## Related Information

[Remote Connection Details \[page 1255\]](#)

[Remote Statement Details \[page 1255\]](#)

[Tile Catalog: Smart Data Access Administration \[page 42\]](#)

## 9.1.7.1 Remote Connection Details

Detailed information about the remote connections active in the database, including when the connection was opened, how many remote statements were executed, and the name of the remote source.

The table below lists the information available for remote connections.

Detail	Description
Connection	Connection ID
SAP HANA cockpit: Adapter	Name of the adapter used for Smart Data Access
Status	Connection status: <ul style="list-style-type: none"><li>• <i>Connected</i>: Connection is active</li><li>• <i>Disconnected</i>: Connection has been closed</li></ul>
SAP HANA studio: 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

## 9.1.7.2 Remote Statement Details

Detailed information about the remote statements executed in the database, including when the query was started, how long the query took, and the number of records returned. The full SQL text of the SQL statements is also shown.

The table below lists the 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)	

Detail	Description
Status	Query execution status: <ul style="list-style-type: none"> <li>Analyzing: Query is being analyzed by the query optimizer</li> <li>Optimizing: Query is being optimized by the query optimizer</li> <li>Executing: Query is running</li> <li>Closed: Query has completed</li> <li>Failed: Query execution failed</li> </ul>
Rows	Number of rows returned in the query result
Remote Source Name	Name of the remote data source
SAP HANA cockpit: User	User who executed the statement
SAP HANA cockpit: Transaction	Transaction ID

## 9.1.8 Remote Connection Pooling

A connection pool can be used to enable multi-threaded execution to be scaled out. Connection pooling can be used only for read-only remote sources (DML\_MODE=readonly).

### **i** Note

Provided no external updates (by third parties) occur on the remote source, then 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.

## Remote Connection Pool Configuration Parameter

To configure remote connection pooling, you use the parameter `default_connections_pool_max_size` in the `smart_data_access` section of the `indexserver.ini` file:

```
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 multi-threading 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.

## 9.1.9 Smart Data Access System Parameters

Configuration parameters for smart data access available in the `indexserver.ini` file in the `smart_data_access` section.

Parameter Name	Type	Length	Values	Default Value	Description	Hidden
enable_remote_source_capability	BOOLEAN		TRUE = any query in the remote source dialect can be sent for remote execution; FALSE = only projections are sent for remote execution	TRUE	Specifies the complexity of queries to be sent to the remote sources.	NO
virtual_table_format	VARCHAR	16	ROW = row based; COLUMN = column based; AUTO = let the optimizer choose	ROW	Forces optimizer to use between column or row-based operators.	NO
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	Specifies the preferred order of semi-join execution strategies.	NO
semi_join_max_in_elements	INTEGER		Positive integer value	1024	Specifies maximum number of values in the IN clause for semi-join usage.	NO
semi_join_min_temp_table_cardinality	INTEGER		Positive integer value		Minimum number of values to be inserted in a semi-join temp table.	YES
semi_join_max_temp_table_cardinality	INTEGER		Positive integer value	16384	Maximum number of values to be inserted in a semi-join temp table.	NO
semi_join_virtual_table_threshold	TINYINT		Positive integer value		Minimum number of estimated rows for fact sub-plan, to be considered for semi-join reduction.	YES

Parameter Name	Type	Length	Values	Default Value	Description	Hidden
semi_join_reduction_factor	TINYINT		Positive integer value		The estimated percentage reduction required for an attribute to be considered for semi-join reduction.	YES

## 9.2 SAP HANA Hadoop Integration

You can combine the in-memory processing power of SAP HANA with Hadoop's ability to store and process huge amounts of data, regardless of structure.

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.

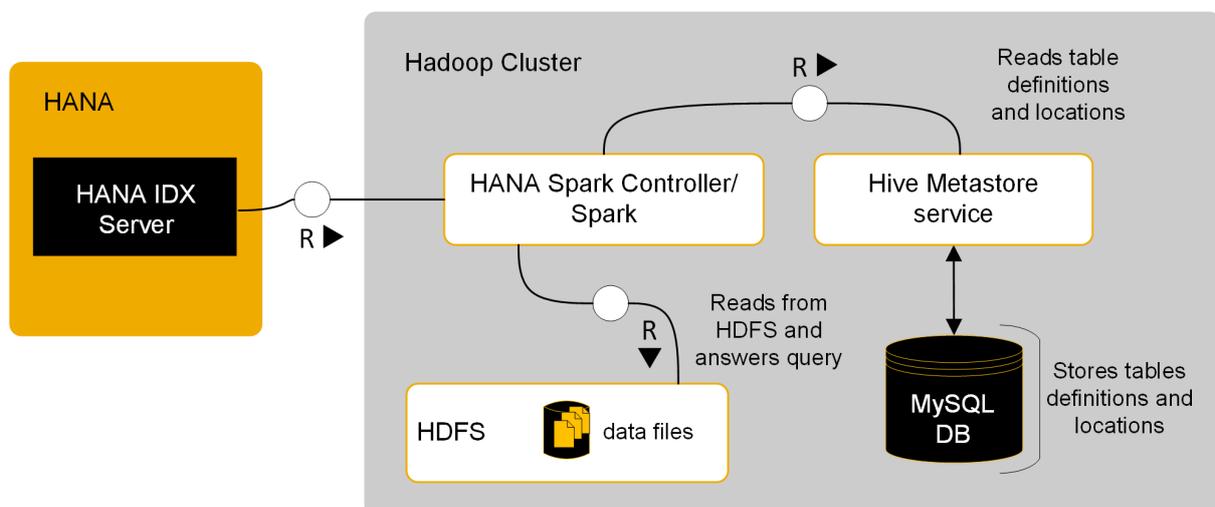
SAP HANA Hadoop integration is designed for users who may want to start using SAP HANA with their Hadoop ecosystem. This document assumes you have a Hadoop cluster installed.

### Integrating SAP HANA and Hadoop

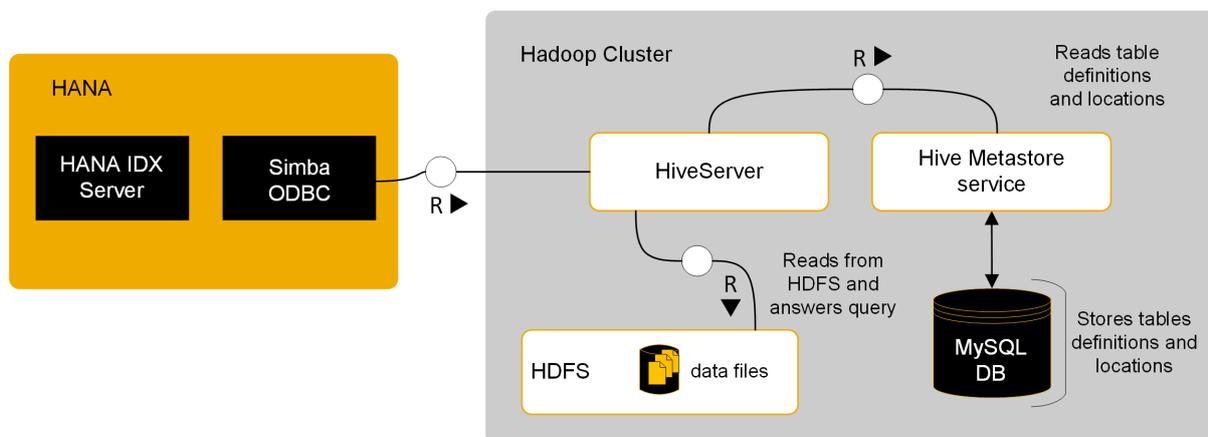
There are two methods available to set up communication between SAP HANA and your Hadoop system:

- (Recommended) SAP HANA Spark controller – See [SAP HANA Hadoop Integration](#).
- Hive ODBC driver – See [Using the Simba ODBC Driver to Connect to Hive \[page 1265\]](#).

This diagram illustrates the role that Spark controller and Spark SQL adapter play in the communication between SAP HANA and a Hadoop cluster:



This diagram illustrates the role that the ODBC drivers play in the communication between SAP HANA and a Hadoop cluster:



## 9.2.1 Hadoop Integration Platform Support

SAP HANA/Hadoop integration is supported on Intel-based hardware platforms and IBM Power Systems, with some exceptions.

- SAP HANA Hadoop controller (Intel-only)
- Hive ODBC driver (Intel-only)
- Any connected Hadoop cluster (or one running Vora) must be installed on a non-Power platform.

## 9.2.2 SAP HANA Spark Controller

SAP HANA Spark controller supports SAP HANA in-memory access to data in the Hadoop cluster HDFS data files.

Spark controller allows SAP HANA to access Hadoop data through the SQL interface and primarily works with Spark SQL to connect to an existing Hive metastore.

Spark controller is assembled, installed, and configured on the Hadoop cluster. YARN and Spark Assembly JAR are used to connect to the HDFS system, with YARN as the resource management layer for the Hadoop ecosystem.

On the SAP HANA side, Spark SQL adapter is a plug-in for the smart data access framework, providing access to Spark controller, and moderating query execution and data transfer.

For installation and configuration instructions, see [SAP HANA Hadoop Integration](#)

---

## 9.2.3 SAP HANA Ambari Integration

The Apache Ambari integration with SAP HANA cockpit allows you to enter the Ambari Web URL in the cockpit and access Hadoop cluster monitoring functionality using Ambari Web UI.

After entering the Ambari Web URL, you can navigate to the Apache Ambari website and monitor Hadoop clusters. You can also use Ambari to set up Spark controller.

### 9.2.3.1 Adding Ambari URL to SAP HANA Cockpit

Add Ambari to the SAP HANA cockpit.

#### Context

After going to the Ambari Web URL, you can navigate to the Apache Ambari website and monitor Hadoop clusters.

#### Procedure

1. Import the cockpit delivery unit package (`HANA_HADOOP_AMBR.tgz`) into SAP HANA studio.
2. Assign these roles to all users requiring access the web application site (requires SAP HANA System Administrator role):
  - `com.sap.hana.hadoop.cockpit.ambari.data::Administrator`
  - `sap.hana.uis.db::SITE_DESIGNER`
  - `sap.hana.uis.db::SITE_USER`
3. In the Systems view, right-click on the system name and select ► *Configuration and Monitoring* ► *Open SAP HANA Cockpit* ► to launch the SAP HANA cockpit.
4. Log in to the cockpit using the SAP HANA username and password.
5. Select *Hadoop Cluster* on the home page.

If the *Hadoop Cluster* tile is not available, select *Tile Catalog* from the menu and add the *Hadoop Cluster* tile to a desired group.
6. For each cluster, provide a Hadoop cluster name and an Ambari URL (for example, `http://my.ambari.server.url:8080`).
7. Select a Hadoop cluster and click on *Go* to navigate to the Ambari website.

---

## 9.2.4 Data Aging with Hadoop

SAP HANA Data Warehousing Foundation Data Lifecycle Manager supports bidirectional data relocation, allowing you to relocate data in native SAP HANA use cases from SAP HANA persistency to storage locations including Hadoop (Spark SQL), and from Hadoop (Spark SQL) to SAP HANA.

You can model aging rules on SAP HANA tables to relocate "aged" or less frequently used data from SAP HANA tables in native SAP HANA applications.

### Related Information

[SAP HANA Data Warehousing Foundation](#)

## 9.2.5 SAP HANA Vora

SAP HANA Vora provides an in-memory processing engine that runs on a Hadoop cluster and Spark execution framework. Able to scale to thousands of nodes, it is designed for use in large distributed clusters and for handling big data.

The SAP HANA Vora solution is built on the Hadoop ecosystem, which provides a collection of components that support distributed processing of large data sets across a cluster of machines. Hadoop allows both structured as well as complex, unstructured data to be stored, accessed, and analyzed across the cluster.

### Related Information

[SAP HANA Vora](#)

## 9.2.6 Virtual Functions

Create a remote source and add a MapReduce program to SAP HANA before creating virtual functions.

## 9.2.6.1 Create the Remote Data Source

Create the remote data source by running a SQL statement in SAP HANA Studio.

### Context

Use the full qualified domain name of your host. You add the name by running command `hostname -f`.

### Procedure

Run the following statement in the SQL console of SAP HANA studio:

```
CREATE REMOTE SOURCE HADOOP_SOURCE
ADAPTER "hadoop"
CONFIGURATION 'webhdfs_url=http://<full_qualified_domain_name>:
50070;webhcat_url=http://<full_qualified_domain_name>:50111'
WITH CREDENTIAL TYPE 'PASSWORD'
USING 'user=hdfs;password=hdfs';
```

The remote source should appear under [Provisioning](#) > [Smart Data Access](#) > [Remote Source](#).

## 9.2.6.2 Creating a Virtual Function

You can create virtual functions in order for applications to make use of the Hadoop MapReduce jobs and virtual user-defined functions.

### Prerequisites

A remote source is required to create a virtual function. If you are creating a virtual function that uses a MapReduce package, also add it to SAP HANA before creating the virtual function.

### Context

A virtual user-defined function (UDF) provides an abstraction for Hadoop MapReduce jobs. It has no function body and can be used in normal SQL statements.

## Procedure

1. In SAP HANA studio, go to the *SAP HANA Development* perspective and then open the *Project Explorer* view.
2. In the *Project Explorer* view, right-click the shared project folder where you want to create the new virtual function and choose **New > Other... > SAP HANA > Database Development > Hadoop Virtual Function** in the context-sensitive pop-up menu. Select *Next* to select the parent folder and specify the file name for your Hadoop virtual function.
3. Edit the following Hadoop virtual function creation template and save it:

```
VIRTUAL FUNCTION <valid schema>.<function name>()
RETURNS TABLE <return table type>
PACKAGE <hadoop mrjobs archive schema>.<hadoop mrjobs archive name>
CONFIGURATION '<remote proc properties>'
AT <hadoop remote source name>
```

4. Activate the Hadoop virtual function by right-clicking the function name in the *Project Explorer* and selecting **Team > Activate** in the context-sensitive menu.
5. Invoke the virtual function remote source by highlighting the corresponding `SELECT` statement in *SQL Console*, and executing it.
6. Open the *Result* tab in the *SQL Console* to see the job results.

## Results

You have queried a Hadoop file using a custom MapReduce job from SAP HANA using simple SQL.

### Example

Example template:

```
VIRTUAL FUNCTION "SYSTEM"."hadoop.mrjobs.demos::DEMO_VF" ()
RETURNS TABLE ( WORD NVARCHAR(60), COUNT integer)
PACKAGE SYSTEM.WORD_COUNT
CONFIGURATION 'enable_caching=true;mapred_jobchain=[{"mapred_input":"/apps/
hive/warehouse/dflo.db/
region/","mapred_mapper":"com.sap.hana.hadoop.samples.WordMapper","mapred_reduc
er":"com.sap.hana.hadoop.samples.WordReducer"}]'
AT "hadoop.mrjobs.demos:: DEMO_SRC"
```

## Related Information

[Adding a MapReduce Program to SAP HANA \[page 1264\]](#)

## 9.2.6.3 Adding a MapReduce Program to SAP HANA

If you have a Java MapReduce (MR) job, you can push it to the SAP HANA repository as a `.hdbmrjob` repository file. It will include all JAR files for the created MapReduce job. This is a prerequisite step for applications to make use of the Hadoop MapReduce jobs and virtual user-defined functions.

### Prerequisites

You have created a Java project in the *SAP HANA Development* perspective of the SAP HANA studio. For more information about creating a project, see *Create a Project for SAP HANA XS* in the *SAP HANA Developer Guide*.

### Context

SAP HANA Hadoop Controller facilitates MapReduce functionality. It is an adapter that can be installed as a delivery unit in HANA XS engine, and can be pushed to Hadoop. After it is installed, assign the `SAP_HANA_sap.hana.xs.lm.roles::Administrator` role to your SAP HANA user then start the SAP HANA Application Lifecycle Manager to import it as a delivery unit. Downloaded the controller from SAP Marketplace.

### Procedure

1. In SAP HANA studio, go to the *SAP HANA Development* perspective., and open the *Project Explorer* view.
2. In the *Project Explorer* view, right-click the shared project folder where you want to create the new virtual function, and choose **New > Other... > SAP HANA > Database Development > Hadoop MR Jobs Archive** in the context-sensitive menu. Select *Next* to select the parent folder and specify the file name for your Hadoop MapReduce jobs archive. Then select *Next*.
3. Select the Java project folder holding the source of the Hadoop MR jobs program. An archive (JAR file), packing in all the dependencies needed to run the program, is created and associated with the shared project folder.
4. Activate the Hadoop MR jobs archive by right-clicking the archive in the *Project Explorer* and selecting **Team > Activate** in the context-sensitive menu.

### Results

During activation of the design-time object `.hdbmrjob`, a corresponding runtime object, is created and stored in the catalog.

## Next Steps

Now you can submit your MapReduce job files to a Hadoop cluster from a SQL statement. You can also define a virtual user-defined function where you specify the Hadoop MapReduce job name stored as a catalog object.

## Related Information

[Creating a Virtual Function \[page 1262\]](#)

### 9.2.7 Using the Simba ODBC Driver to Connect to Hive

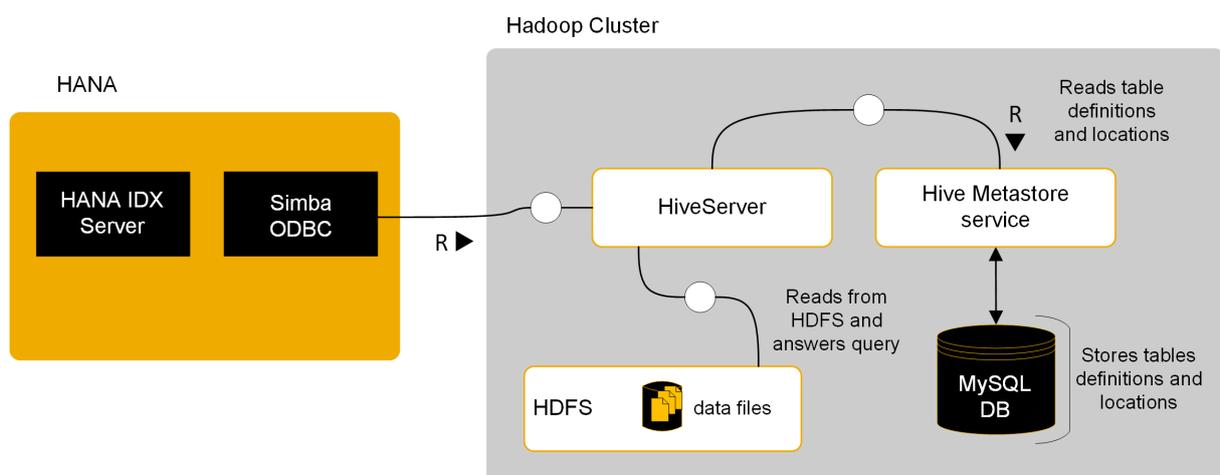
The Simba Hive ODBC Driver is a connector to Apache Hive, a SQL-oriented query language that provides a quick and easy way to work with data stored in HDFS on a Hadoop cluster. Hive leverages Hadoop's processing power and storage capacity.

Installing the ODBC driver allows you to federate data and combine cold data stored in Hadoop with warm data stored in SAP HANA. Using Hive, you can run one SQL query and combine this data.

#### **i** Note

The recommended method for Hadoop and SAP HANA communication is to use Spark controller, which offers performance improvements and additional integration with SAP HANA Vora.

This diagram illustrates the role that the ODBC drivers play in the communication between SAP HANA and a Hadoop cluster:



#### **i** Note

HIVE ODBC Driver is supported on Intel-based hardware platforms only.

## 9.2.7.1 Hive ODBC Overview Task

This section describes an overview of how to set up the communication between SAP HANA and a Hadoop system using the Hive ODBC driver.

1. Set up the Hive ODBC driver. See [Hive ODBC Driver Setup \[page 1266\]](#).
2. Add your Hadoop system as a remote source. See [Create the Remote Data Source \[page 1262\]](#).

After creating a remote source, add a MapReduce program (optional) and access data from Hadoop by creating a virtual function. See:

- [Adding a MapReduce Program to SAP HANA \[page 1264\]](#)
- [Creating a Virtual Function \[page 1262\]](#)

## 9.2.7.2 Hive ODBC Driver Setup

How to set up the Hive ODBC Driver.

### Prerequisites

Make sure the Hive driver is installed on the machine with the SAP HANA instance. You should see it in `/opt/simba/hiveodbc/lib/64/libsimbahiveodbc64.so`

### Procedure

1. Log on to `<sid>adm` using your SAP HANA administration user.
2. Stop SAP HANA.
3. Copy `/opt/simba/hiveodbc/Setup/simba.hiveodbc.ini` to your home directory as `.simba.hiveodbc.ini`:

```
$> cp /opt/simba/hiveodbc/Setup/simba.hiveodbc.ini ~/.simba.hiveodbc.ini
```

4. Edit the `.simba.hiveodbc.ini` file with the following:
  - If there is a line with `DriverManagerEncoding=UTF-32`, change the value to `UTF-16`.
  - Make sure the line `ErrorMessagePath=/opt/simba/hiveodbc/ErrorMessage` exist.
  - Comment out the line: `ODBCInstLib=libiodbcint.so`.
  - Uncomment the line: `ODBCInstLib=libodbcinst.so`.
5. Open or create your `.odbc.ini` file in your home directory and add the following:

```
[hive1]
Driver=/opt/simba/hiveodbc/lib/64/libsimbahiveodbc64.so
Host=server.com
HIVE Port=10000
```

---

`Host` is the machine running Hive. `HIVE Port` is the port running Hive (the default is 10000).

6. Add the following environment variables to the `$HOME/.customer.sh` file.
  - a. Add `/opt/simba/hiveodbc/lib/64/` to your `LD_LIBRARY_PATH`. For example:

```
$> export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/opt/simba/hiveodbc/lib/64
```

- b. Set the `ODBCINI` environment variable to your `.odbc.ini` file location:

```
$> export ODBCINI=$HOME/.odbc.ini
```

7. Restart the Hadoop server and verify that you can connect to Hive through the ODBC driver using the DSN name. For example:

```
isql $>isql -v hive1
```

8. Start SAP HANA.

### 9.2.7.3 Downloading Supporting Libraries

Third-party libraries are needed to support the processing of date or time, and comma-separated files.

#### Procedure

1. If it does not already exist, create the following Hadoop Distributed File System (HDFS) directory in Hadoop:

```
/sap/hana/mapred/lib
```

2. If they do not already exist, add the following libraries:

- `joda-time-2.3.jar`

Download from <http://mvnrepository.com/artifact/joda-time/joda-time/2.3>.

- `solr-commons-csv-3.5.0.jar`

Download from <http://mvnrepository.com/artifact/org.apache.solr/solr-commons-csv/3.5.0>.

### 9.2.7.4 Enable Remote Caching

When using the Hive ODBC driver to connect SAP HANA and Hadoop, you can enable remote caches in Hive for queries on low-velocity data, which allows you to use materialized data for the repetitive execution of the same query.

When SAP HANA dispatches a virtual table query to Hive, it involves a series of map and reduce job executions. Completing a could take hours, depending on the data size in Hadoop and the current cluster capacity. In most cases, the data in the Hadoop cluster is not frequently updated, and successive execution of map and reduce jobs might result in same queries. Using remote caching with Hadoop through the Hive interface allows you to use the cached remote data set so you need not wait for map reduce to be executed. There is no performance improvement the first time you run a statement because of the time it takes to run the map reduce job and the

sorting of data in the table. The next time you run the same query, you are accessing the materialized data, therefore reducing the execution time for the job.

Use this feature for Hive tables with low data velocity (which are not frequently updated).

This behavior is controlled by using a hint to instruct the optimizer to use remote caching. For example:

### Example

You have created a virtual table `hive_activity_log` and then fetch all erroneous entries for plant 001:

```
select * from hive_activity_log where incident_type = 'ERROR' and plant = '001'  
with hint (USE_REMOTE_CACHE)
```

When you use the hint `USE_REMOTE_CACHE`, this result set is materialized in Hive and subsequent queries are served from the materialized view.

## Configuration Parameters

The following configuration parameters are used for remote caching and are stored in the `indexserver.ini` file in the `smart_data_access` section:

Parameter	Description
<code>enable_remote_cache ( 'true'   'false' )</code>	A global switch to enable or disable remote caching for federated queries. This parameter only supports Hive sources. The <code>USE_REMOTE_CACHE</code> hint parameter is ignored when this parameter is disabled.
<code>remote_cache_validity = 3600 (seconds)</code>	Defines how long the remote cache remains valid. By default, the cache is retained for 1 hour.

## 9.3 Data Replication and Transformation

Overview of replicating and transforming data.

Smart data integration is a set of functionality provided by several components you can use to retrieve data from an external system, transform it, and persist it in SAP HANA database tables. Smart data integration features and tools addressed in this section include:

- The Replication Editor in the SAP HANA Web-based Development Workbench is used to create real time or batch replication scenarios for moving data into SAP HANA studio.
- The Transformation Editor in the SAP HANA Web-based Development Workbench is used to create flowgraphs for transforming data, such as filtering, sorting and joining data into tables.
- Likewise, the application function modeler in SAP HANA Studio is also used to create flowgraphs for transforming data similar to SAP HANA Web-based Development Workbench.

## Related Information

[Replicating Data \[page 1269\]](#)

[Transforming Data Using SAP HANA Web-based Development Workbench \[page 1280\]](#)

[Transforming Data Using Application Function Modeler \[page 1291\]](#)

[Node Reference \[page 1322\]](#)

### 9.3.1 Replicating Data

Replicate data from several objects in a remote source to tables in SAP HANA using the Replication Editor in SAP HANA Web-based Development Workbench.

To replicate data from objects in a remote source into tables in SAP HANA, you must configure the replication process by creating an `.hdbreptask` file, which opens a file specific to the Replication Editor.

Before using the Replication Editor, you must have the proper rights to use the editor. See your system administrator to assign appropriate permissions. You must also have the run-time objects set up as described in the "SAP HANA Web-based Development Workbench: Catalog" chapter of the *SAP HANA Developer Guide*.

#### **i** Note

The Web-based Editor tool is available on the SAP HANA XS web server at the following URL: `http://<WebServerHost>:80<SAPHANAInstance>/sap/hana/ide/editor`.

After the `.hdbreptask` has been configured, activate it to generate a stored procedure, a remote subscription, one or more virtual tables for objects that you want to replicate, and target tables. The remote subscription is only created when the *Initial load only* option is cleared. When the stored procedure is called, an initial load is run. When real time is enabled, then subsequent changes are automatically distributed.

DDL changes to source tables that are associated with a replication task will be propagated to SAP HANA so that the same changes will be applied to the SAP HANA target tables.

See the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality Administration Guide* for information about monitoring and processing remote subscriptions for realtime replication tasks.

## 9.3.1.1 Create a Replication Task

A replication task retrieves data from one or more objects in a single remote source and populates one or more tables in SAP HANA.

### Prerequisites

Before using the Replication Editor, you must have the proper rights to use the editor. For example, you must have the ALTER Object privilege on the remote source where you'll be searching. See your system administrator to assign appropriate permissions.

### Context

Columns that you do not map as inputs to a replication task or flowgraph will not be sent over the network to be processed by SAP HANA. Excluding columns as inputs can improve performance, for example if they contain large object types. You can also choose to enhance security by excluding columns that, for example, include sensitive data such as passwords.

### Procedure

1. Highlight a package from the content pane and right-click. Choose ► *File* ► *New* ► *Replication Task* .
2. Enter a file name and then click *Create*.
3. In *Remote Source*, select the source data location from the drop-down list.
4. In *Target Schema*, select the schema for the target table.
5. In *Virtual Table Schema*, select the schema for the virtual table.
6. Select whether to *Use Package Prefix* for the virtual and/or target tables. For example, if your virtual table name is customer\_demo, and you enable the *Virtual Table* option, the output would be "VT\_customer\_demo".
7. Set *Drop target table if exists* with one of the following:

Target table	Outcome
Existing	When this option is selected and there is an existing target table, the target table is deleted and recreated.
New	When this option is selected, a new target table is created.

8. (Optional) In the *Virtual Table Prefix* option, enter some identifying letters or numbers to help you label the virtual table. You might want a prefix to identify where the data came from or the type of information that it contains.
9. To include one or more tables in the replication task, click *Add Objects*.

10. In the *Select Remote Source* window, you can browse to or search for the object(s) as follows. Note that in the *Importable* column, a zero means that it is not importable and a one means that it can be imported. You can use Shift-click or Ctrl-click to select multiple objects.

- To browse for the object, expand the nodes as necessary and select the object(s).
- To search for an object:
  - Click *Create Dictionary* to build a searchable dictionary of objects from the source.

### **i** Note

You only need to create the dictionary the first time you search for an object. It is automatically available after the first search.

- Enter filter criteria for *Display Name*, *Unique Name*, or *Object Description* that *Contains*, *Equals*, *Starts with*, or *Ends with* characters you enter.  
For example, to filter by name, enter the first few characters of the object name to display the objects that begin with those characters. The *Case sensitive* restriction is optional. To add additional criteria to further filter the list, click the plus sign and enter the additional parameter(s).
  - (Optional) The bottom of this interface includes a time stamp for when the dictionary was last updated. You can also refresh or clear the dictionary here.
  - Select the object(s) to add.
11. (Optional). Enter a prefix in the *Target name prefix* option. For example, you might want the prefix to be ADDR\_ if the output table contains address data. The rest of the table name is the same as the remote object name. You can change the entire name on the main editing page, if necessary.
12. (Optional) Configure the *Replication Behavior* for the table. You can choose to perform a combination of initial load, realtime replication, and table-level structure replication.
13. With the desired object(s) selected, click *OK*.
14. (Optional) Click the *Filter* tab to enter SQL statements to further limit the rows being replicated using the SQL syntax of a WHERE clause. Only records that meet the criteria of the filter are replicated.
15. Click *Save*.

## Related Information

[Add a Target Column \[page 1272\]](#)

[Edit a Target Column \[page 1273\]](#)

[Delete a Target Column \[page 1273\]](#)

[Replication Behavior Options for Objects in Replication Tasks \[page 1274\]](#)

[Load Behavior Options for Target Tables \[page 1275\]](#)

[Activate and Execute a Replication Task \[page 1278\]](#)

## 9.3.1.2 Add a Target Column

Add a column in a replication task.

### Procedure

1. From the Replication Editor, in the *Target Columns* tab, click *Add*.
2. Choose whether to create a column or to include a column from a remote object.
  - *From remote object*: Browse to a source and table and choose the column you replicated in the virtual table. Select whether the column *is part of the primary key*, and then click *OK*.
    1. Select the column name.
    2. Select if this column *is part of the primary key*.
    3. Click *OK*.
    4. Rename the column.
    5. Enter the projection.
  - *From scratch*: Complete the following steps to create a column. Then, you can enter some SQL statements in the *Filter* tab to set the value of the target column during replication. Any of the SAP HANA SQL functions can be used. See the *SAP Hana SQL and System Views Reference*.
    1. Enter the *Name* of the column.
    2. Select the *Data Type*. For example, varchar, decimal and so on.
    3. Enter the number of characters allowed in the column.
    4. Enter the *Projection* (the mapped name) of the column.

#### **i** Note

The projection can be any one of the following:

- column (enter the name of the source column in double quotes, for example, "APJ\_SALES")
- string literal (enter the string as a value in single quotes, for example 'ERPCLNT800')
- SQL expression (for example, "firstname" + "lastname")

5. Select *is nullable* if the value can be empty.
6. Select *is part of the primary key* if the data in the column will uniquely identify each record in a table.
7. Click *OK*.

### Related Information

[Create a replication task \[page 1270\]](#)

---

### 9.3.1.3 Edit a Target Column

Modify the column to correct the data or to make it more accurate or useful.

#### Context

For example, if you were using a Social Security number as a part of a primary key, and you need to stop using it for the primary key, you can edit the column to unselect the option. To edit a column:

#### Procedure

1. Select the column.
2. Click *Edit*.
3. Change the data type, length, projection, nullable, and/or primary key options.
4. Click *OK*.

### 9.3.1.4 Delete a Target Column

Remove a column so that it is no longer used in the flowgraph.

#### Procedure

1. Select the column.
2. Click *Delete*.
3. Confirm your deletion, and then click *OK*.

## 9.3.1.5 Replication Behavior Options for Objects in Replication Tasks

For replication tasks, you can select different options that control initial data loading, real-time time replication, and object structure replication.

### Context

Replication behavior is managed at the target object level in the Replication Editor. For example, you can choose to enable realtime replication for one target table, and perform only an initial data load for another table within the same replication task.

Set the replication behavior for a target object by selecting the object in the Replication Editor and choosing the desired behavior from *Replication Behaviors*. You can set the behavior for multiple target objects by selecting the objects and choosing *Set Replication Behavior*.

### Procedure

1. Select the replication task in the Workbench Editor.
2. Select the Remote Object to edit.
3. In the *Details* pane, select the *Target Columns* tab.
4. From the *Replication Behavior* drop-down menu, select one of the following options:

Behavior	Description
<b>Initial load only</b>	Performs a one-time data load without any realtime replication
<b>Initial + realtime</b>	Performs the initial data load and enables realtime replication
<b>Realtime</b>	Enables realtime replication without performing an initial data load
<b>No data transfer</b>	Replicates only the object structure without transferring any data
<b>Initial + realtime with structure</b>	Performs the initial data load, enables realtime replication, and tracks object-level changes
<b>Realtime only with structure</b>	Enables realtime replication and tracks object-level changes without performing an initial data load

5. Save the replication task.

## 9.3.1.5.1 Support for Schema Change Replication

When you choose a replication behavior that includes the object structure, the types of schema changes that can be replicated depend on the target adapter.

Replication support per adapter

Schema Change	IBM DB2 Log Reader	Microsoft SQL Server Log Reader	Oracle Log Reader	SAP ECC	SAP HANA	Teradata
Add columns	Yes	Yes	Yes	Yes	Yes	Yes
Drop columns	Yes	Yes	Yes	Yes	Yes	Yes
Alter column data type		Yes		Yes**		
Rename column		Yes		Yes**		
Rename table		Yes		Yes**		

### Restriction

Adapters not mentioned do not support schema change replication.

Schema change replication for *Alter column data type*, *Rename column*, and *Rename table* is supported on SAP ECC only on a Microsoft SQL Server database. These schema changes are not supported during replication on IBM DB2 or Oracle databases.

## 9.3.1.6 Load Behavior Options for Targets in Replication Tasks

For real-time replication tasks, you can select different options that enable one-to-one replication, actuals tables, or change log tables as targets.

### Context

Simple replication of a source table to a target table results in a copy of the source (same row count, same columns). However because the table replication process also includes information on what row has changed and when, you can add these change types and change times to the target table.

For example, in simple replication, deleted rows do not display in the target table. To display the rows that were deleted, you can select the *Actuals Table* option that functions as UPSERT when loading the target. This option adds two columns CHANGE\_TYPE and CHANGE\_TIME to the target table. The deleted rows display with a CHANGE\_TYPE of D.

You can also choose to display all changes to the target (INSERT functionality) which provides a change log table. Every changed row is inserted into the target table including the change types, change time, and a sequence indicator for multiple operations that were committed in the same transaction.

Column	Description																
CHANGE_TYPE	<p>Displays the type of row change in the source:</p> <table border="1"> <tbody> <tr> <td>I</td> <td>INSERT</td> </tr> <tr> <td>B</td> <td>UPDATE (Before image)</td> </tr> <tr> <td>U</td> <td>UPDATE (After image)</td> </tr> <tr> <td>D</td> <td>DELETE</td> </tr> <tr> <td>A</td> <td>UPSERT</td> </tr> <tr> <td>R</td> <td>REPLACE</td> </tr> <tr> <td>T</td> <td>TRUNCATE</td> </tr> <tr> <td>X</td> <td>EXTERMINATE_ROW</td> </tr> </tbody> </table>	I	INSERT	B	UPDATE (Before image)	U	UPDATE (After image)	D	DELETE	A	UPSERT	R	REPLACE	T	TRUNCATE	X	EXTERMINATE_ROW
I	INSERT																
B	UPDATE (Before image)																
U	UPDATE (After image)																
D	DELETE																
A	UPSERT																
R	REPLACE																
T	TRUNCATE																
X	EXTERMINATE_ROW																
CHANGE_TIME	Displays the time stamp of when the row was committed. All changes committed within the same transaction will have the same CHANGE_TIME.																
CHANGE_SEQUENCE	Displays a value that indicates the order of operations for changes that were committed in the same transaction.																

## Procedure

1. Select the replication task in the Workbench Editor.
2. Select the Remote Object to edit.
3. In the *Details* pane, select the *Load Behavior* tab.
4. From the *Load Behavior* drop-down menu, select one of the following options:
  - *Replicate*: Replicates changes in the source one-to-one in the target.
  - *Replicate with logical delete*: UPSERTS rows and includes CHANGE\_TYPE and CHANGE\_TIME columns in the target.
  - *Preserve all*: INSERTS all rows and includes CHANGE\_TYPE, CHANGE\_TIME, and CHANGE\_SEQUENCE columns in the target.
5. (Optional) You can rename the column names.
6. Save the replication task.

### Example

Consider the following changes made to the LinItem table for sales order 100:

Operation	Time stamp	Description
Insert	08:01	Add new line item 3 worth \$60
Insert	08:02	Add new line item 4 worth \$40
Delete	08:02	Delete line item 1
Commit	08:03	Save the changes to the order

The target tables would display as follows.

*Replication Table:*

Order	Line	Material	Amount
100	2	Bolt	200
100	3	Nut	60
100	4	Spacer	40

*Actuals Table:*

Order	Line	Material	Amount	CHANGE_TYPE	CHANGE_TIME
100	1	Screw	200	D	2015-04-23 08:04
100	2	Bolt	200	I	2015-04-23 08:04
100	3	Nut	60	I	2015-04-23 08:04
100	4	Spacer	40	I	2015-04-23 08:04

*Change Log Table:*

Order	Line	Material	Amount	CHANGE_TYP E	CHANGE_TIM E	CHANGE_SE- QUENCE
100	1	Screw	200	I	2015-04-23 07:40	23
100	2	Bolt	200	I	2015-04-23 07:40	24
100	3	Nut	60	I	2015-04-23 08:04	50
100	4	Spacer	40	I	2015-04-23 08:04	51
100	1	Screw	200	D	2015-04-23 08:04	52

## Related Information

[Load Behavior Options for Targets in Flowgraphs \[page 1288\]](#)

### 9.3.1.7 Activate and Execute a Replication Task

Activation generates the run time objects necessary for data movement from one or many source tables to one or more target tables.

#### Context

The replication task creates the following run time objects.

- Virtual table(s): Generated in the specified virtual table schema. You can display the contents of the virtual table in SAP HANA studio.
- Remote subscription(s): Generated in the schema selected for the virtual table. This is only generated when the *Initial load only* option is not selected.
- Task(s): Generated in the same schema as the target table.
- View(s): Generated in the same schema as the virtual table.
- Target table(s): Populated with the content after execution.
- Procedure: Generated in the schema of the target table, the procedure performs three functions.
  1. Sets the remote subscription to the Queue status.

#### **i** Note

The remote subscription is only created when *Initial load only* is unselected.

2. Calls Start Task to perform the initial load of the data.
3. Sets the remote subscription to the Distribute status. Any changes, additions or deletions made to the source data during the initial load are updated in the target system. Any changes to the source data thereafter are updated real time to the target.

#### **i** Note

The remote subscription is only created when *Initial load only* is unselected.

#### Procedure

1. After the replication task is configured, click [Save](#) to activate.
2. Go to the Catalog view and navigate to the stored procedure you just created.

### **i** Note

You can access the Catalog view on the SAP HANA XS Web server at the following URL `http://<WebServerHost>:80<SAPHanaInstance>/sap/hana/xs/ide/catalog`. Choose one of the following options to activate the replication task.

- Right-click the stored procedure, and then select *Invoke Procedure*.
- To call the stored procedure, use the following SQL script:

```
CALL
"<schema_name>". "<package_name>::<target_table_name>".START_REPLICATION
```

The replication begins. You can right-click and select *Open Contents* to view the data in the target table in the Catalog view.

### **i** Note

If the replication task takes longer than 300 seconds to process, you might receive an error about the XMLHttpRequest failing. You can correct this issue by increasing the maximum run time option in the `xsengin.ini` file. Follow these steps:

1. Login to SAP HANA studio as a SYSTEM user.
2. In the Systems view, right-click the name of your SAP HANA server, and then choose **► Configuration and Monitoring ► Open Administration ►**.
3. Click the *Configuration* tab.
4. Select `xsengine.ini`.
5. Expand *httpserver*.
6. Click *Add parameter*.
7. In the *Assign Values to* option, select *System*, and then *Next*.
8. In the *Key* option, enter `max_request_runtime` and then enter a value. For example, you might want to enter 1200. The value is in seconds.
9. Click *Finish* and then close the *Configuration* tab and execute the replication task again.

## Results

You can use SAP HANA Cockpit to monitor the results.

## Related Information

[SAP HANA SQL and System Views Reference \(HTML\)](#)

---

## 9.3.2 Transforming Data Using SAP HANA Web-based Development Workbench

Use SAP HANA Web-based Development Workbench for replicating and transforming data.

Use this tool to develop applications in a Web browser without having to install any development tools. This is a quick alternative to using SAP HANA Studio's application function modeler. It simplifies development by providing many convenient functions such as replicating data into HANA, and transforming that data so that you are using the records and tables necessary for your business. When creating a flowgraph with the nodes in the General palette, you will create a procedure that you can call after activation.

Columns that you do not map as inputs to a replication task or flowgraph will not be sent over the network to be processed by SAP HANA. Excluding columns as inputs can improve performance, for example if they contain large object types. You can also choose to enhance security by excluding columns that, for example, include sensitive data such as passwords.

### 9.3.2.1 Add a Variable to the Flowgraph

Create variables to simplify the process of activating a flowgraph.

#### Context

When you create variable, you can use them in nodes that accept them such as the Filter and Join nodes. For example, in a Filter node, you might want to process only those records for a certain country, such as Spain. You can create a variable for each country in the flowgraph properties. Then you can call the variable in the filter by surrounding the variable name with \$\$\$. For example,

```
"COUNTRY" = $$$Spain$$$
```

#### Procedure

1. In the *Properties*, click *Variables*.
2. Click *Add*.
3. Enter values for the variable.

Option	Description
Name	The name of the variable. For example, "Florida". When using the variable in other nodes, surround the variable name with two dollar signs. For example, in the Filter node when you output Florida data, you would use  <pre>"STATE" = \$\$Florida\$\$</pre>
Kind	Select one of the following options.  <i>Expression</i> : Use in nodes where the expression editor is located. This includes filters and attribute values.  <i>Scalar Parameter</i> : Use with scalar parameters such as R script procedures. There must be one scalarParam for each variable in this Variables tab.  <i>Task</i> : Use when creating a flowgraph (task) level variable. You can use this variable during flowgraph partitioning.
Type	The type of data contained in the column, for example, Nvarchar, Decimal, Date, and so on. Required when using scalarParam.
Length	The number of characters allowed in the column. Required when using scalarParam.
Scale	The number of digits to the right of the decimal point. This is used when the data type is a decimal. Required when using scalarParam.
Nullable	Indicates whether the column can be null.
Default	Enter a value to use when the criteria is not met in the node. For example, when using the Filtering node to look for customers in Germany, and you might set the default to Berlin if the country is not specified.

## Results

Then when you activate the flowgraph, you can specify the output by calling the variable(s) in the function. For example,

```
START TASK "<schema_name>"."<package_name>:::<flowgraph_name>" (country =>
  'US', state => 'NY');
```

## Related Information

[SAP HANA SQL and System Views Reference \(PDF\)](#)

[SAP HANA SQL and System Views Reference \(HTML\)](#)

## 9.3.2.2 Partitioning Data in the Flowgraph

Partitioning data can be helpful when you are initially loading a large data set, because it can improve performance. This topic contains general information about partitioning, and specific information about single-level partitioning.

### Context

Data partitioning is used to separate large data sets into smaller sets based on a set of defined criteria. These partitions can be run in serial or in parallel. Some common reasons for partitioning include:

- You receive “out of memory” errors when you load the data.
- You have reached the limit for the maximum number of rows within the column store.
- You want the performance to be faster.

You can partition data for the flowgraph using the following:

- Virtual tables
- Physical tables
- Calculation views
- SQL views
- Input type

You can partition data in two ways: at the task (flowgraph) level, and at the Data Source node level. Partitioning at the task level is useful when your input data has several million rows or more. Currently, SAP HANA has a limitation of processing more than two billion rows. Partitioning your data at the task level will likely reduce the load to less than two billion rows per partition. Typically, you only see a benefit of using task level partitioning with extremely large data sets. You can set the number of parallel partitions that are processed simultaneously. The transformation and loading to the target is done per partition. When you partition at the task level, you must select one data source in the flowgraph.

If your input data is smaller, then it might be better to use the partitioning options in the Data Source node, although partitioning can only be done on virtual tables. When partitioning in the Data Source node, only the input data is partitioned. All of the partitions are run in parallel; you cannot change the number of parallel partitions.

If you partition data in both the Data Source node and at the task level, the task level partition settings take precedence, and the specified Data Source node settings are ignored. If you have multiple Data Source nodes defined with partitioning, only the data source selected in the task partitioning is impacted during runtime. All other Data Source nodes that are partitioned within the task will be processed with their individual partitioning settings.

Partitioning can have an impact on your data results. When you partition at the task level, you must select one data source in the flowgraph. Because you can have multiple data sources within a flowgraph and only one can be partitioned at the task level, there may be slower performance when using a Join node, or different data results when using the Match node.

#### **i** Note

The Match node may not be available to you depending on your application version or license agreement.

When using the Join node, one data source is partitioned, and the other Data Source nodes are used in their entirety. If those sources have a significant amount of data, the Join node may act as a bottleneck because it has to wait for the other Data Source node data to load before being able to join the contents with the task level partitioned data source.

Regarding the Match node results, the Match node will only find duplicates from within a single partition of data. However, if the partitioning is set on data that makes good break keys, then the Match node results will not be an issue.

To create a single-level task partition:

## Procedure

1. Click the *Properties* icon.
2. Set the *Runtime Behavior Type* to *Batch Task*.
3. (Optional) In the *Variables* tab, define any task variables that you want to use in the flowgraph. See [Add a Variable to the Flowgraph \[page 1280\]](#).
4. Click the *Partitions* tab and choose *Task*.
5. In the *Input Source* list, you will see the supported data sources in your flowgraph. You can only set partitions for one of these sources. Typically, this is the source with the largest set of data.
6. In the *Column* option, choose the column that you want to use as the base of your partitioning. You will set your partitions based on the data in this column. If you will use the Column partition type, then you do not need to set this option.
7. Set the *Number of Parallel Partitions*. For example, if you have five partitions, and you set the Number of Parallel Partitions to 2, then partitions 1 and 2 are run together. If partition 2 completes before partition one, then partition 3 starts running, and so on. In general, the more partitions that you run in parallel, the faster the data is loaded. However, if there are memory issues, then data may not load because there are too many parallel partitions set. When this option is set to 1, then partitioning is run in a series beginning with the first partition. When that partition is finished, the second partition is started, and so on.

You can set this option with a number or a variable. If you have defined a task variable in the *Variables* tab, you can enter it here with two dollar signs surrounding it. For example, `$$variable_name$$`. Use only a task variable with a positive integer value. See the [Add a Variable to the Flowgraph \[page 1280\]](#) for adding task variables.

8. Select the Partition Type that you want to use.

Option	Description
<b>Column</b>	The data must be partitioned at the table level before it can be used in the flowgraph (columnTable). The IDs are automatically assigned based on the table settings. The Attribute option does not need to be set because the columnTable partitioning information will be used.
<b>List</b>	The data is divided into sets based on a list of values in a column. For example, if you want to partition France, Germany and the United Kingdom, you would enter 'FR', 'DE', 'GB' with single quotes around each value.
<b>Range</b>	The data is divided into sets based on a range of data in the column. You need to enter only the ending value in the range. For example, the range of values can be

Option	Description
	<p>300. If this is an integer value, it includes all records from 0-299. (For a non-integer value, it includes values less than 300.) You could have a second partition that has the ending value of 900, which includes all records from 300-899.</p> <p>String value results are in alphabetical order. When using string values, make sure the data is surrounded by single quotation marks, for example, 'NM' for the state of New Mexico in the United States. Note that in the string value, the partition would not include the value 'NM'. In this example, it would include all US states from AK (Alaska) through NJ (New Jersey).</p>

9. Click the + icon.
10. Enter a *Partition Name*. For example, if you are partitioning on the country names, you might enter the location of the country, like Western Europe.
11. Enter the *Values* that you want included in the partition. For example, you might enter 'FR', 'PT', 'ES', 'GB' for the countries of France, Portugal, Spain, and United Kingdom. Click *OK*. Repeat the previous three steps to add more partitions. You must have a minimum of two partitions defined.

#### Note

The last partition must be created with a blank value so that it captures any remaining records. Then, none of your input data is lost.

12. When you have finished adding partitions, click *OK* to return to the Flowgraph Editor.

### Example

#### Column Partitioning Example

Let's say that you have a table of Canadian census data that is already partitioned by a Region column name.

Partitioning at the column table level:

Partition ID	Value
1	Alberta
2	British Columbia
3	Manitoba
4	New Brunswick
5	Newfoundland and Labrador
6	Northwest Territories
7	Nova Scotia
8	Nunavut
9	Ontario
10	Prince Edward Island
11	Quebec
12	Saskatchewan

Partition ID	Value
13	Yukon

Partitioning at the table and task level might increase performance. For this example, you want to partition the data based on the provinces that have largest population, in this case, Quebec, Ontario, British Columbia and Alberta. You might set your partitions like this:

Partitioning: Task	Number of parallel partitions: 2
Input Source: Canada_Census_Data	Partition type: Column

Partition Name	Value
Quebec	11
Ontario	9
British Columbia	2
Alberta	1
Other	<blank>

Because there are two partitions, Quebec and Ontario partitions are started together, when one of the partitions is finished, then British Columbia is started, and so on. Those records that contain '1', '2', '9' and '11' values are put in their respective partitions. All other records are placed in the Other partition.

#### List Partitioning Example

Let's say that you have a table with European customers. You have hundreds of thousands of customers in Spain, France and Germany, and tens of thousands of customers in Belgium, Netherlands and Denmark. You might set your partitions like this:

Partitioning: Task	Number of parallel partitions: 2
Input Source: Euro_Data	Partition type: List
Column: Country	

Partition Name	Value
Spain	'ES'
France	'FR'
Germany	'DE'
Other	<blank>

Because there are two partitions, Spain and France partitions are started together. When one of the partitions is finished, then Germany started, followed by the Other partition. Those records that contain 'ES', 'FR', and 'DE' values are put in their respective partitions. All other records are placed in the Other partition.

#### Range Partitioning Example

Let's say that you have a large amount of data for the state of New York, and you want to load your data based on the postcode range. Because the majority of the data is in New York City, you've decided to split those postcodes into 3 partitions.

Partitioning: Task	Number of parallel partitions: 1
Input Source: New_York_Data	Partition type: Range
Column: Postcode	

Partition Name	Value
NYC1	10100
NYC2	10200
NYC3	12288
Other_NYC	<blank>

Because the number of parallel partitions is 1, then the data is loaded serially. You only need to specify the ending value for the range. Any numeric values prior to that are included in the partition. For example, NYC1 lists the end value of 10100. This partition includes all numbers from 00000-10099. NYC2 contains postcodes from 10100-10199, and NYC3 contains postcodes from 10200-12287. All records that are not specified in the first 3 partitions are placed in the Other\_NY partition.

### 9.3.2.2.1 Multi-level Partitioning

Multi-level partitioning further divides the partitioned data at the task (flowgraph) level.

#### Context

Multi-level partitioning can take a considerable amount of planning. Data can be lost when the partitions are not correctly set. Therefore, it is recommended that multi-level partitioning be set by advanced users.

To create a multi-level partition:

#### Procedure

1. Follow the steps in the topic [Partitioning Data in the Flowgraph \[page 1282\]](#).
2. Select *Use multi-level partitioning*.
3. Under *Partition Levels*, click the + icon.
4. Select the *Partition Type* that you want to use.

Option	Description
Column	The data must be partitioned at the table level before it can be used in the flowgraph (columnTable). The IDs are automatically assigned based on the table settings. If you want to use column partitioning, you must use it as your first partition.

Option	Description
	You can only use column partitioning once. Any additional partitions must be List and/or Range types.
<b>List</b>	The data is divided into sets based on a list of values in a column. For example, if you want to partition France, Germany and the United Kingdom, you would enter 'FR', 'DE', 'GB' with single quotes around each value.
<b>Range</b>	The data is divided into sets based on a range of data in the column. For example, the range of values can be 0,300 and includes all records with those values. You could have a second partition that has the values 301,900. When using string values, make sure the data is surrounded in single quotation marks, for example 'CA', 'NY'.

5. Select the column that you want to base the first level partition on.
6. Under the *Partitions*, click the + icon to define the sub-partitions based on the levels above. You will enter a value for each of the levels.
7. Enter a *Partition Name*.
8. Enter a *Value* for each of the partitions.
9. Repeat the previous three steps to add more partitions. You must have a minimum of two partitions defined.

### Note

It is recommended that you have set values for each of the partitions so that all of the records in the input source are placed in a partition. However, if you are unsure that you have captured all of the values in that column, create a partition with a blank value so that it captures any remaining records. Then, none of your input data will be lost.

10. 10. When you have all of the partitions set, click *OK* to return to the Flowgraph Editor.

### Example

Let's say that you have a popular product (ID #22456) that is sold basic and premium levels. It is very popular in parts of North America (Canada, Mexico, and United States) and Asia (China, Japan, and Republic of Korea), so you want to partition the data based on product, premium/basic and country.

Partitioning: Task	Number of parallel partitions: 3
Input Source: Product_Sales	Use multi-level partitioning

Partition levels:

Level	Type	Column
L1	List	Product_ID
L2	List	Level
L3	List	Country

Partitions:

Level	L1	L2	L3
Basic_Asia	'22456'	'Basic'	'CN', 'JP', 'RK'
Premium_Asia	'22456'	'Premium'	'CN', 'JP', 'RK'
Basic_NA	'22456'	'Basic'	'CA', 'US', 'MX'
Premium_NA	'22456'	'Premium'	'CA', 'US', 'MX'
Other	<blank>	<blank>	<blank>

Because there are five partitions and three parallel partitions set, Basic\_Asia, Premium\_Asia and Basic\_NA are started together. When one of the partitions completes, then Premium\_NA begins, and then Other will begin when the next partition is completed. The specified data is placed into the appropriate partitions. All other records are placed in the Other partition.

### 9.3.2.3 Load Behavior Options for Targets in Flowgraphs

For flowgraphs, you can select options that enable different target-loading behaviors and include columns that display the time and type of change made in the source.

#### Context

Simple replication of a source table to a target table results in a copy of the source (same row count, same columns). However, because this process also includes information on what row has changed and when, you can add these change types and change times to the target table.

For example, in simple replication, deleted rows do not display in the target table. To display the rows that were deleted, you can select UPSERT when loading the target. The deleted rows display with a change type of D.

You could also choose to display all changes to the target (INSERT functionality), which provides a change log table. Every changed row would be inserted into the target table and you can include columns that display the change types and change times.

Column	Description	
CHANGE TYPE	Displays the type of row change in the source:	
	I	INSERT
	B	UPDATE (Before image)
	U	UPDATE (After image)
	D	DELETE
	A	UPSERT
	R	REPLACE
	T	TRUNCATE
	X	EXTERMINATE_ROW
CHANGE TIME	Displays the time stamp of when the row was committed. All changes committed within the same transaction will have the same change time.	

## Procedure

1. As a prerequisite for INSERT operations, in the SQL Console, create a sequence.

```
CREATE SEQUENCE "DPUSER"."SEQ_QA_EMP_HISTORY" START WITH 1 INCREMENT BY 1
MAXVALUE 4611686018427387903 MINVALUE 1 ;
SELECT "DPUSER"."SEQ_QA_EMP_HISTORY".NEXTVAL FROM DUMMY;
```

2. For an existing target table, add columns to the table for storing change types, change times, and change sequence numbers.
3. Add or open a flowgraph in the Workbench Editor.
4. Open the target editor.
5. In the *Node Details* pane on the *General* tab, select a **Writer Type** (insert or upsert).
6. On the *Settings* tab:
  - a. Select a *Key Generation Attribute*.
  - b. Select a *Sequence Name*.
  - c. Select a *Sequence Schema*.
  - d. Select the previously configured *Change Time Column Name*.

If the target is a template table, you can select an existing column or type a new name to create a new target table.

- e. Select the previously configured *Change Type Column Name*.

If the target is a template table, you can select an existing column or type a new name to create a new target table.

7. Save the flowgraph.
8. Activate the flowgraph.

## Related Information

[Data Sink Options \[page 1327\]](#)

[Load Behavior Options for Targets in Replication Tasks \[page 1275\]](#)

### 9.3.2.4 Activate and Execute a Flowgraph

After your flowgraph is created and configured, activate it to create the run-time objects.

#### Context

Activation creates the run-time objects based on the options set in the flowgraph.

#### Procedure

1. From the Project Explorer, right-click on the `.hdbflowgraph` that you created.
2. Choose **► Team ► Activate ►**.  
The run time objects are created.
3. Choose one of the following:
  - If you configured the flowgraph for initial load only, use the following SQL to run the generated task:

```
START TASK "<schema_name>"."<package_name>::<flowgraph_name>"
```

#### **i** Note

You can also specify a variable when running Start Task. For example, if you have a Filter node set to output records for a specific country, you can enter it in a similar way to the following.

```
START TASK "<schema_name>"."<package_name>::<flowgraph_name>" (country => 'Spain');
```

- If you configured the flowgraph for real time, use the following SQL script to execute the generated initialization procedure:

```
CALL "<package_name>::<flowgraph_name>_SP"
```

- If you configured the flowgraph for real time and want to pass a variable value, use the following script to execute the generated initialization procedure:

```
CALL "<package_name>::<flowgraph_name>_SP"('"'Spain'"')
```

For more information about Start Task and calling a table type, see the “Start Task” topic.

## Related Information

[SAP HANA SQL and System Views Reference \(HTML\)](#)

### 9.3.3 Transforming Data Using SAP HANA Application Function Modeler

Overview of SAP HANA application function modeler.

The SAP HANA application function modeler is the default editor for flowgraphs. A flowgraph is a development object. It is stored in a project and has extension `.hdbflowgraph`. By default, the activation of a flowgraph generates a procedure in the catalog.

#### **i** Note

If the optional additional cost SAP HANA smart data integration and SAP HANA smart data quality component is available, a flowgraph can be configured to generate a task plan run-time object instead of a procedure.

#### **i** Note

Columns that you do not map as inputs to a replication task or flowgraph will not be sent over the network to be processed by SAP HANA. Excluding columns as inputs can improve performance, for example if they contain large object types. You can also choose to enhance security by excluding columns that, for example, include sensitive data such as passwords.

A flowgraph models a data flow that can contain:

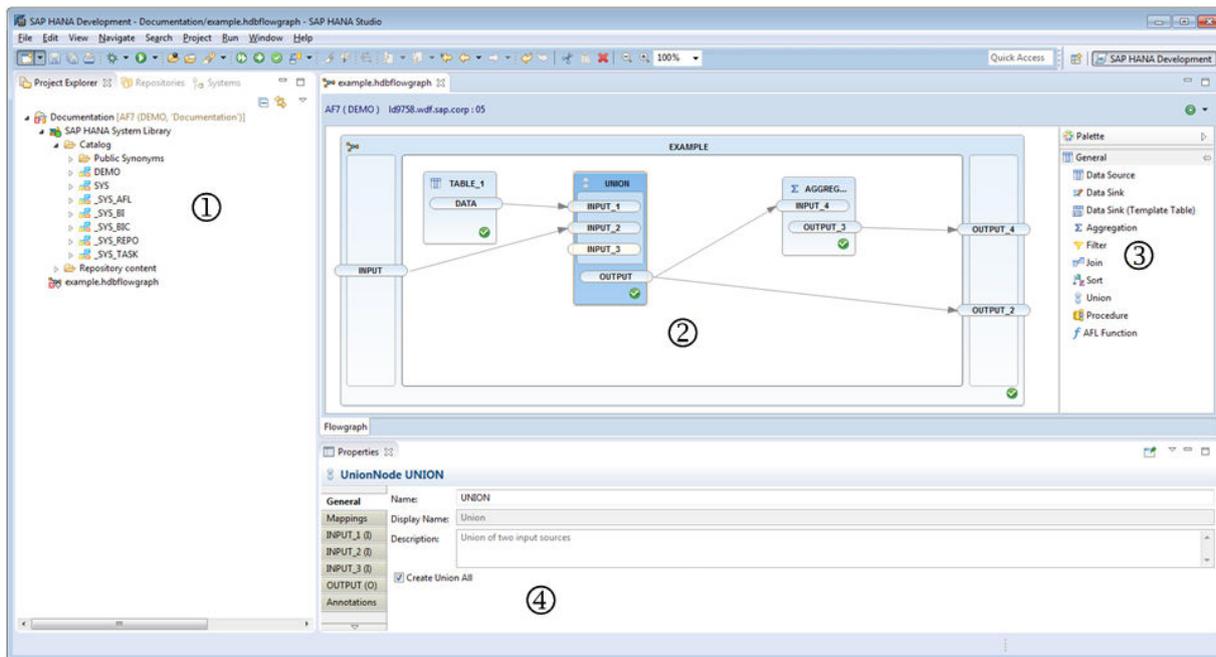
- tables, views, and procedures from the catalog
- relational operators such as projection, filter, union, and join
- functions from Application Function Libraries (AFL) installed on your system
- attribute view and calculation view development objects

In addition, the application function modeler provides support for some optional, additional cost components of the SAP HANA Platform such as:

- the Business Function Library
- the Predictive Analysis Library
- R Scripts

- Data Provisioning operators
- the generation of task plans

The application function modeler is part of the *SAP HANA Development* perspective and utilizes the following components.



Components used by the SAP HANA application function modeler

	Area	Description
1	<i>Project Explorer</i> view	The <i>Project Explorer</i> is used as a source of objects that can be added to the <i>Editing Area</i> .
2	<i>Editing Area</i>	In the <i>Editing Area</i> , the flowgraph is modeled. Elements are added to the flowgraph by dragging objects from the <i>Project Explorer</i> or node templates from the <i>Node Palette</i> to the <i>Editing Area</i> . There, they can be selected and edited via the context button pad and the context menu. The <i>Editing Area</i> supports standard editing operations like copy, paste, and delete, as well as moving elements by drag and drop. The properties of selected flowgraph elements can be edited in the <i>Properties</i> view.
3	<i>Node Palette</i>	The <i>Node Palette</i> lists the node templates available to the application function modeler. These node templates can be added to the flowgraph by dragging them to the <i>Editing Area</i> . In case an optional, additional cost component of the SAP HANA Platform is detected by the application function modeler, an additional compartment with node templates for its functions is automatically added to the <i>Node Palette</i> .

	Area	Description
4	<i>Properties</i> view	The <i>Properties</i> view shows the property details of the selected flowgraph element.

#### ➔ Tip

You can open the *SAP HANA Development* perspective by choosing **Window > Open Perspective > SAP HANA Development**, the *Properties* view by choosing **Window > Show View > Properties**, and the *Project Explorer* views by choosing **Window > Show View > Project Explorer**.

## Related Information

[Add a Variable to the Flowgraph \[page 1280\]](#)

### 9.3.3.1 Converting deprecated AFL Models (AFLPMML objects)

Convert a deprecated AFL Model development object that was created by a previous version of the SAP HANA application function modeler into a flowgraph.

#### Context

AFL Models are development objects with the extension `.aflpmml` that were created with a previous version of the SAP HANA application function modeler. They are deprecated in SAP HANA SPS09.

Compared to the complex data flows with various operators modeled by a flowgraph, an AFL Model object is restricted to model a single function from the Application Function Library together with the data sources and data sinks that are connected to this function.

An AFL Model can still be activated. However, since AFL Models are deprecated, it can no longer be directly edited with the application function modeler. Instead, the AFL Model first has to be converted to a flowgraph. Then this flowgraph can be edited with the application function modeler. For backward compatibility, the edited flowgraph can be re-converted to an AFL Model. This requires all changes to the flowgraph to be compatible with the restrictions of AFL Models.

#### Procedure

1. In the *Project Explorer* view right-click on the AFL Model that you want to convert to a flowgraph, and then choose *Convert to Flowgraph* in the context-sensitive menu.

---

The application function modeler creates a new flowgraph with the same prefix and the `.hdbflowgraph` extension. A dialog appears that lets you delete the AFL Model and its corresponding generated procedure. Afterward, you can edit the new flowgraph with the application function modeler.

#### **i** Note

If you choose not to delete the converted application function modeler Model and try to activate a flowgraph, you get an error stating that there already exists an active catalog object with the same name (the new object tries to generate the same runtime object). You need to either delete or rename one of the two objects and activate the modification as well.

#### **i** Note

A flowgraph cannot be activated on a SAP HANA SPS08 system.

2. (Optional) Convert a flowgraph to a AFL Model. In the *Project Explorer* view right-click on the flowgraph that you want to convert to an AFL Model, and then choose *Convert to AFLPMML* in the context-sensitive menu. The application function modeler creates a new AFL Model with the same prefix and the `.aflpmm1` extension.

#### **i** Note

AFL Model objects are deprecated. This conversion is available for backward compatibility. Most features of a flowgraph are not supported by the AFLPMML format.

## 9.3.3.2 Setting up the SAP HANA Application Function Modeler

Configure your system to use the SAP HANA Application Function Modeler.

Before modeling flowgraphs with the SAP HANA Application Function Modeler (AFM), make sure that the following system requirements are satisfied and that the following database access rights are granted to the respective database users.

### System Requirements

The AFM has the following system requirements.

- You have installed the current version of SAP HANA.
- You have installed the Application Function Libraries (AFLs) that you want to use. For more information, see the section *Installing or Updating SAP HANA Components* in the *SAP HANA Server Installation and Update Guide*.
- You have enabled the Script Server in your SAP HANA instance. See SAP Note 1650957 for more information.

## Privileges for the database user `_SYS_REPO`

The database user `_SYS_REPO` has to be granted the following object privileges:

- SELECT object privileges for objects that are used as data sources,
- INSERT object privileges for objects that are used as data sinks,
- INSERT and DELETE object privileges for objects that are used as data sinks with truncation.

### **i** Note

Granting access rights to the user `_SYS_REPO` may constitute a security risk, so make sure that you understand the privileges you grant to database users. For more information, see the *SAP HANA Security Guide*.

## Privileges for the database user of the AFM

You have to be granted the MODELING role.

You have to be granted the EXECUTE privilege for the object `SYS.REPOSITORY_REST`.

You have to be granted the following package privileges:

- `repo.read` package privileges on your repository package
- `repo.activate_native_objects` package privileges on your repository package
- `repo.edit_native_objects` package privileges on your repository package
- `repo.maintain_native_packages` package privileges on your repository package

In addition, you have to be granted the following object privileges to the target schema of the flowgraph activation (default: `_SYS_BIC`):

- CREATE ANY
- ALTER
- DROP
- EXECUTE
- SELECT
- INSERT
- UPDATE

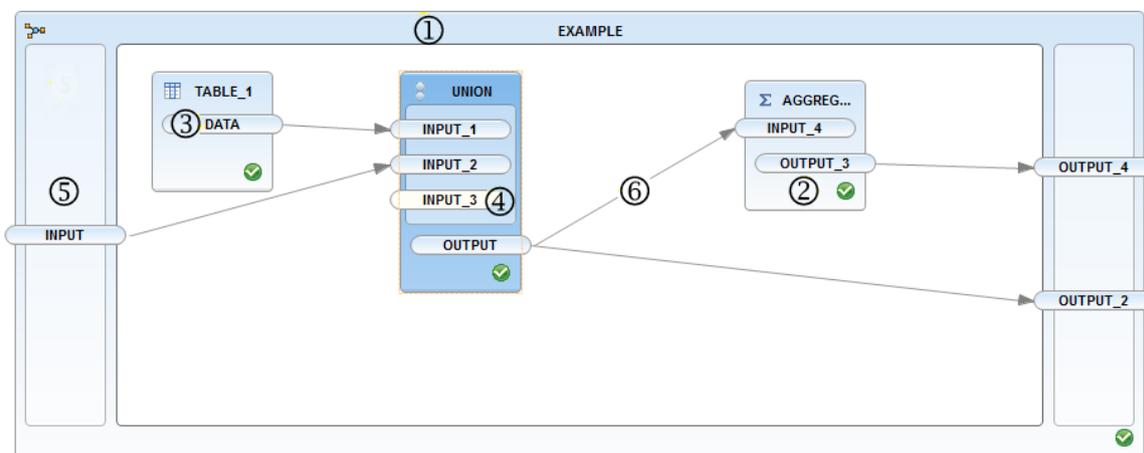
### **i** Note

Granting access rights to the user `_SYS_REPO` may constitute a security risk, so make sure that you understand the privileges you grant to database users. For more information, see the *SAP HANA Security Guide*.

### 9.3.3.3 Flowgraphs

This is an overview of all flowgraph elements.

A flowgraph consists of several flowgraph elements that are depicted in the *Editing Area*. Every flowgraph element has a collection of properties that are displayed in the *Properties* view.



Flowgraph Elements

	Element	Description
1	Flowgraph container	The flowgraph container represents the operator defined by the flowgraph. Every flowgraph has exactly one flowgraph container. This flowgraph container has a name which has to differ from all other elements of the flowgraph. The flowgraph container can have several anchors. They represent the inputs and outputs of the operator defined by the flowgraph. The central free area of the flowgraph container is its canvas. All nodes of the flowgraph are contained in this canvas. The validation decorator in the right bottom corner of the flowgraph container indicates whether the flowgraph is configured correctly.
2	Node	Nodes are the functional elements in a flowgraph. There are several different types of nodes which represent data sources, data sinks, and operators. A node has a name which has to be unique in the flowgraph. Like the flowgraph container, a node can have several anchors. They represent the inputs and outputs of the node. The validation decorator in the right bottom corner of a node indicates whether the node is configured correctly.

	Element	Description
3	Anchor	An anchor represents an input or an output of the flowgraph container or of a node. Every anchor has a kind and a signature which define the input or output it represents. For input anchors, the supported kinds are Table, Column, and Scalar. For output anchors, the only supported kind is Table. Anchors of the Column kind are considered to be tables with a single column. Anchors of the Scalar kind are considered to be tables with a single column and a single row. This way, every anchor defines the table type of the input or output it represents.
4	Fixed content anchor	A fixed content anchor is an input anchor for which the fixed content flag is set in the properties. It is displayed in white color (in contrast to the light-blue colored standard anchors). A fixed content anchor cannot be the target of a connection. Instead, there is a table embedded in the flowgraph that is associated to the fixed content anchor. The table is displayed in the <i>Fixed Content</i> tab of the <i>Properties</i> view of the anchor.
5	Anchor region	The flowgraph container and some nodes (for example, the Join node and the Union node) can have a variable number of input and output anchors. This is represented by an anchor region for the corresponding set of anchors. Anchors can be added to or removed from the anchor region. They can also be reordered in the anchor region.
6	Connection	Connections represent the directed flow of data from a source to a target. The source and the target of a connection are anchors. The connection defines a table mapping between the table types defined by its source and target. The source of a connection is either an input anchor of the flowgraph container or an output anchor of a node. The target of a connection is either an output anchor of the flowgraph container or an input anchor of a node. An anchor can be the source of several connections. It can be the target of only one connection. A fixed content anchor cannot be the target of a connection.

## Validation

There is a validation decorator in the right bottom corner of the flowgraph container and of each node. This decorator indicates if the complete flowgraph or the respective node is configured correctly. The details of a validation error are displayed by mouse-over on the validation decorator and in the *Problems* view.

### ➔ Tip

You can open the Problems view by choosing **Window > Show View > Problems** in the main menu of the HANA Studio.

### ⚠ Caution

A flowgraph with validation errors will fail to activate.

## Annotations

Annotations are nested key-value pairs that can be added to the flowgraph container and to nodes. The AFM uses annotations to store certain properties of the flowgraph such as custom palette information. An AFM user

---

can store arbitrary meta data in the annotations. When the flowgraph is activated, all annotations are exposed in a table with the name extension `.META` in the flowgraph target schema. This way, they can be consumed at runtime.

There are two main reasons for the user of the AFM to create annotations. The first reason is to add comments and documentation to the flowgraph. The second reason is to pass meta data about the flowgraph and its nodes to an application consuming the runtime procedure generated by the activation. In this case the application has to be specifically designed to process the meta data. Although this is a rather specific and uncommon use-case, it is a very versatile approach that utilizes flowgraphs to configure the analytic functionality of an application.

## Related Information

[Using the Mapping Editor \[page 1311\]](#)

[Using the Annotation Editor \[page 1313\]](#)

[Customizing the Node Palette \[page 1319\]](#)

### 9.3.3.4 Modeling a flowgraph

Model a flowgraph starting with its creation and concluding with the execution of the generated procedure.

#### Context

The SAP HANA Application Modeler (AFM) supports standard graphical editing operations like move, copy, paste, and delete on the elements of a flowgraph. Detailed properties of these elements are edited in the *Properties* view. After editing and saving a flowgraph, it can be activated by the AFM and the generated procedure can then be executed via the AFM.

If the flowgraph container has input anchors, the procedure has corresponding free inputs. It then cannot be executed directly. In this case, data sources have to be bound to the free inputs in order to execute the runtime object. The AFM provides a wizard for this.

#### Procedure

1. Create a new flowgraph or open an existing flowgraph in the *Project Explorer* view.  
The flowgraph is opened in the *Editing Area* of the AFM.
2. Edit the flowgraph container.
3. Add and edit nodes.
4. Add and edit anchors.

5. Add and edit connections.  
The validation decorators in the bottom right corners of the flowgraph container and the nodes indicate whether the flowgraph is valid.

**i Note**

A flowgraph must be valid to be activated.

6. Save the flowgraph. Select **File > Save** in the HANA Studio main menu.
7. Activate the flowgraph. In the *Project Explorer* view, right-click the flowgraph object and choose **Team > Activate...** from the context-sensitive menu.  
A new procedure is generated in the target schema which is specified in the properties of the flowgraph container.

**i Note**

The generated procedure has inputs that correspond to the input anchors of the flowgraph container. To activate this procedure, these inputs have to be specified.

8. Select the black downward triangle next to the *Execute* button  in the top right corner of the AFM.  
A context menu appears. It shows the options *Execute in SQL Editor* and *Open in SQL Editor* as well as the option *Execute and Explore* for every output of the flowgraph. In addition, the context menu shows the option *Edit Input Bindings*.
9. (Optional) If the flowgraph has input anchors, choose the option *Edit Input Bindings*.  
A wizard appears that allows you to bind all inputs of the flowgraph to data sources in the catalog.
10. Choose one of the options *Execute in SQL Editor*, *Open in SQL Editor*, or *Execute and Explore* for one of the outputs of the flowgraph.  
The behavior of the AFM depends on the execution mode.

Execution mode	Behavior
<i>Open in SQL Editor</i>	Opens a SQL console containing the SQL code to execute the runtime object.
<i>Execute in SQL Editor</i>	Opens a SQL console containing the SQL code to execute the runtime object and runs this SQL code.
<i>Execute and Explore</i>	Executes the runtime object and opens the <i>Data Explorer</i> view for the chosen output of the flowgraph.

11. Close the flowgraph. Select **File > Close** in the HANA Studio main menu.

## 9.3.3.4.1 Creating a Flowgraph

You can use the SAP HANA AFM to create procedures using PAL functions. The first step is to create a Flowgraph (.hdbflowgraph) file.

### Procedure

1. In the *SAP HANA Development* perspective, open your existing SAP HANA project.
2. In the *Project Explorer* view, right-click on the project name, and choose ► *New* ► *Other* ►.  
The *New* wizard will appear.
3. From the Wizard list, expand *SAP HANA* and then *Database Development*, select *Flowgraph Model*, and choose *Next*.
4. Enter or select the parent folder, and enter the file name.

#### Note

You only need to enter the base name. The system automatically adds the extension .hdbflowgraph to it.

5. Choose *Finish*.

The new .hdbflowgraph file will appear in the *Project Explorer* view with the  icon.

#### Note

If your workspace was created in SAP HANA SPS05 or earlier, you will also see a .diagram file and a .af1model file. Opening these two files may cause errors. To avoid it, you need to manually activate the Diagram-af1model filter as follows:

1. In the *Project Explorer* view, click the drop-down arrow (View Menu) in the upper-right corner and choose *Customize View*.
2. In the *Available Customizations* dialog box, select *Diagram-af1model Filter* and choose *OK*.

6. If your project has not yet been shared, right-click on the project name, select ► *Team* ► *Share* ► *Project* ►.  
Make settings in the *Share Project* wizard and choose *Finish*.

## 9.3.3.4.2 Editing the flowgraph container

Edit the properties of the flowgraph in the *Properties* view of the flowgraph container.

### Context

Tabs in the Properties view of the flowgraph container.

Tab name	Description	Optional
<i>General</i>	This tab contains the following entries: <ul style="list-style-type: none"><li>• <i>Name</i>: name of the flowgraph container,</li><li>• <i>Display Name</i>; not used,</li><li>• <i>Description</i>; not used,</li><li>• <i>Target Schema</i>: schema in which the runtime object is generated during activation (default: <code>_SYS_BIC</code>),</li><li>• <i>Generator</i>: the type of runtime object to be generated during activation. The option <i>Task</i> is only relevant if the flowgraph uses the additional cost SAP HANA smart data integration and SAP HANA smart data quality optional component,</li><li>• <i>Realtime Behavior</i>: This option is only relevant when the flowgraph uses the additional cost SAP HANA smart data integration and SAP HANA smart data quality optional component and the chosen <i>Generator</i> option is <i>Task</i></li></ul>	No
<i>Variables</i>	This tab is relevant only when the flowgraph uses the SAP HANA smart data integration and SAP HANA smart data quality optional component. For more information, see the "Add a Variable to the Container Node" topic in the <i>Modeling Guide for SAP HANA smart data integration and SAP HANA smart data quality</i> .	Yes
<i>Mappings</i>	The <i>Mapping Editor</i> in this tab is used to remove or re-order input and output anchors and their attributes.	Yes
<i>INPUT (I) / OUTPUT (O)</i>	These tabs correspond to the input and output anchors of the flowgraph container. They have the same names as the respective anchors and the same contents as the <i>All</i> tabs in the <i>Properties</i> views of the anchors.	Yes
<i>Annotations</i>	This tab contains the annotations of the flowgraph container.	No
<i>All</i>	This tab is a summary of all tabs in this view except for the input and output anchor tabs.	No

### Procedure

1. Select the flowgraph container and open the *Properties* view.
2. In the *General* tab, specify the name of the flowgraph container, as well as the target schema, and the generator of the flowgraph.

### **i** Note

The name of the flowgraph container is initially auto-generated from the name of the flowgraph object in the *Project Explorer* view. This name has to be changed if it does not adhere to the naming rules for the flowgraph elements. Names of flowgraph elements may contain only upper-case letters, underscores, and digits and must be unique in the flowgraph.

### **i** Note

You need to be granted the CREATE ANY, ALTER, DROP, EXECUTE, SELECT, INSERT, and UPDATE privileges to the target schema of the flowgraph.

3. In the *Mappings* tab, use the *Mapping Editor* to remove or reorder input and output anchors and their attributes.
4. In the *Annotations* tab, use the *Annotations Editor* to edit the annotations of the flowgraph container.

## Results

The settings made on the flowgraph container determine the type of runtime object generated during activation and the number and signatures of its inputs and outputs.

## Related Information

[Flowgraphs \[page 1296\]](#)

[Using the Mapping Editor \[page 1311\]](#)

[Using the Annotation Editor \[page 1313\]](#)

[Customizing the Node Palette \[page 1319\]](#)

[Add a Variable to the Flowgraph \[page 1280\]](#)

### 9.3.3.4.3 Adding an object from the Project Explorer

Drag and drop an object from the *Project Explorer* view to the *Editing Area*.

## Prerequisites

You have opened a flowgraph in a project that has been shared with a HANA system.

## Context

Nodes are the functional elements in a flowgraph. There are several types of nodes which represent data sources, data sinks, and operators in the flowgraph.

The following database objects are represented by nodes in a flowgraph.

- Development objects in the project:
  - Flowgraphs with no inputs and one output as Data Source nodes
  - Attribute Views as Data Source nodes
  - Calculation Views as Data Source nodes

### **i** Note

Flowgraphs that represent procedures with inputs or with more than one output cannot be directly inserted in other flowgraphs. However, it is possible to add the procedure generated by activating one flowgraph to another flowgraph. This is done via drag and drop from the catalog (see below) or by adding a Procedure node from the *Node Palette*.

- Runtime objects in the catalog:
  - Tables as Data Source nodes, and as Data Sink nodes
  - Views as Data Source nodes
  - Procedures without scalar parameters and INOUT parameters as Procedure nodes
  - Table Types and Tables as anchors

### **➔** Tip

You can also drag a Table Type or a Table to an anchor region to create a new anchor.

## Procedure

- In the Project Explorer, select an object and drag it to the canvas of the flowgraph container.

If the dragged object is a table, a pop-up dialog lets you choose if this table is used as a data source or a data sink in the flowgraph.

A new node is added to the flowgraph. The type of the node matches the selected object in the Project Explorer. The flowgraph container is re-sized so that the new node is contained in the canvas of the flowgraph container.

### **i** Note

You need to be granted SELECT access rights on the schema that contains the object.

### **i** Note

In order to activate the flowgraph, database user `_SYS_REPO` needs to be granted SELECT object privileges for objects that are used as data sources and INSERT object privileges for objects that are used as data sinks.

### Caution

The validation of the SAP HANA Application Function Modeler does not recognize when the signature of an input or output of a table or view has changed. In this case the signature of the respective input or output of the added node is inconsistent with that of the object. Consequently, the flowgraph activation fails.

- In the Project Explorer, select a table type or a table and drag it to an anchor region.

A new anchor with the same signature as the table type or the table is added to the anchor region at the position where the object was dropped.

### Note

You need to be granted SELECT access rights on the object.

### Note

Dragging a table to the anchor region only transfers the signature of the table to the anchor. No reference to the table or its content is stored in the flowgraph. Accordingly, no additional object privileges have to be granted to the database user `_SYS_REPO`.

## Related Information

[Setting up the SAP HANA Application Function Modeler \[page 1294\]](#)

### 9.3.3.4.4 Adding a node from the Node Palette

Drag a node template from the *Node Palette* to the canvas of the flowgraph container in the *Editing Area*.

## Prerequisites

You have opened a flowgraph in a project that has been shared with a HANA system.

### Note

The *Node Palette* is generated according to the functionality provided by the system. If you work in a project that is not shared with a system or the system is offline, the content of the *Node Palette* is restricted to a few basic relational operators. For example, the Data Source node and Data Sink node will be missing the *General* tab.

---

## Context

Nodes are the functional elements in a flowgraph. There are several different types of nodes which represent data sources, data sinks, and operators in the flowgraph.

## Procedure

In the Node Palette, select the entry you want to add and drag it to the canvas area of the flowgraph container.

## Results

A new node is added to the flowgraph. The type of the node matches the selected node template in the [Node Palette](#). The flowgraph container is resized such that the new node is contained in its canvas.

## Related Information

[Setting up the SAP HANA Application Function Modeler \[page 1294\]](#)

### 9.3.3.4.5 Editing a node

Edit the properties of a node.

## Context

The nodes in a flowgraph usually need to be configured. Relational nodes need configurations such as join conditions, filter predicates, and attribute sets for projection. Edit the configuration of a node by selecting the node and navigating to its [Properties](#) view. The selection of tabs and the configuration options in the [Properties](#) view depend on the type of node.

Tabs in the Properties view of a node

Tab name	Description	Optional
<i>General</i>	<p>This tab always contains the following elements:</p> <ul style="list-style-type: none"> <li>• <i>Name</i>: name of the node (editable),</li> <li>• <i>Display Name</i>: name of the node template entry (read-only),</li> <li>• <i>Description</i>: description of the node template entry (read-only).</li> </ul> <p>In addition, this tab contains most configuration options that are specific to the particular node type.</p>	No
<i>Script</i>	<p>This tab is only relevant if an optional additional cost component offers Script node functionality (for example, R Integration).</p>	Yes
<i>Mappings</i>	<p>If the node defines a mapping of its inputs to its outputs or contains an anchor region, this mapping is displayed and can be edited in the <i>Mapping Editor</i>.</p>	Yes
<i>INPUT (I) / OUTPUT (O)</i>	<p>These tabs correspond to the input and output anchors of the node. They have the same names as the respective anchors and the same contents as the <i>All</i> tabs in the <i>Properties</i> views of the anchors.</p>	Yes
<i>Annotations</i>	<p>This tab contains the annotations of the node.</p>	No
<i>All</i>	<p>This tab is a summary of all tabs in this view except for the input and output anchor tabs.</p>	No

## Procedure

1. Select a node or add a new node.
2. Select the name of the node.  
The name field becomes active for editing.

### **i** Note

The name of a node may contain only upper-case letters, underscores, and digits. It must be unique within the flowgraph.

3. In the *Annotations* tab of the *Properties* view, use the *Annotations Editor* to edit the annotations of the node.
4. Edit the remaining properties of the node in the *Properties* view. In particular, specify the type-specific properties of the node in the *General* tab.

## Related Information

[Using the Mapping Editor \[page 1311\]](#)

[Using the Annotation Editor \[page 1313\]](#)

## 9.3.3.4.6 Adding an anchor

Add an anchor to an anchor region of the flowgraph container or a node.

### Context

The flowgraph container and some nodes (for example, the Join node and the Union node) can have a variable number of input or output anchors. In the flowgraph, this is represented by the existence of an anchor region for the corresponding set of anchors. New anchors can be added to the anchor region.

### Procedure

1. Right-click on the anchor region at the position you want to add the new anchor.
2. In the context-sensitive menu, choose *Add Input* or *Add Output* (depending on whether you selected an anchor region for inputs or outputs).

### Results

A new anchor with an empty signature is added to the anchor region at the mouse pointer position where the context menu is opened.

#### **i** Note

Instead of adding an anchor via the context-sensitive menu, you can also copy an existing anchor to an anchor region. This has the advantage that the new anchor has a fully defined signature.

#### **i** Note

Alternatively, you can add a new anchor while creating a connection. In this case the new anchor inherits the signature from the source anchor of the connection.

#### **i** Note

A third option to add an anchor with a predefined signature is by dragging a table or a table type from the catalog to the anchor region.

#### **i** Note

You can also delete an anchor that you added to an anchor region. Some anchor regions have a minimum number of anchors (for example, the anchor regions for the inputs of the Join node and the Union node each have to contain at least two anchors). In this case, if the anchor region contains the minimum number of anchors, then no anchor in the anchor can be deleted.

## Related Information

[Adding an object from the Project Explorer \[page 1302\]](#)

[Editing the flowgraph container \[page 1301\]](#)

[Join \[page 1336\]](#)

[Union \[page 1343\]](#)

### 9.3.3.4.7 Editing an anchor

Change and define input and output table types.

## Context

Anchors define inputs and outputs to the flowgraph container and to nodes.

Tabs in the Properties view of an anchor.

Tab name	Description	Optional
General	This tab contains the following entries: <ul style="list-style-type: none"><li><i>Name</i>: name of the anchor</li><li><i>Kind</i>: kind of the anchor (Table, Column, Scalar).</li></ul>	No
Signature	In this tab, you can use the <a href="#">Table Editor</a> to change the signature of the anchor. Anchors of the kind Scalar or Column are considered to be tables with one column.	No
Fixed Content	This tab exists only for input anchors. While the checkbox <i>Fixed Content</i> is selected, the anchor cannot be the target of a connection. Instead, a table providing the input is stored in the flowgraph with the anchor. The table can be edited using the Table Editor in this tab.	Yes
All	This tab is a summary the other tabs.	No

#### **i** Note

Most anchors have a fixed kind that cannot be changed. Currently, the anchor kinds "Column" and "Scalar" are only supported for input anchors of AFL Function nodes.

#### **i** Note

Many anchors either have a fixed signature or obtain their signature via an automatic table mapping.

## Procedure

1. Select the anchor.
2. Select the name of the anchor and edit it in the direct editing area.  
The name field becomes active for editing.

### **i** Note

The name of an anchor must consist of upper-case letters, underscores, and digits. It must be unique in the flowgraph.

3. Use the *Table Editor* to edit the signature of the anchor in the *Signature* tab of the *Properties* view.
4. Select the *Fixed Content* tab in the *Properties* view.
5. If you want to embed the content of the anchor with the flowgraph, select the checkbox *Fixed Content*.  
If the checkbox *Fixed Content* is selected, the embedded table is shown in the *Fixed Content* tab. Use the *Table Editor* to edit the table.

### **i** Note

For some areas of the Application Function Library the SAP HANA application function modeler provides template AFL Function nodes in separate compartments of the *Node Palette*. These template nodes are preconfigured with fixed signature tables if the respective input is a design-time parameter of the node.

## Related Information

[Flowgraphs \[page 1296\]](#)

[Using the Table Editor \[page 1310\]](#)

[Using the Mapping Editor \[page 1311\]](#)

### 9.3.3.4.8 Creating a connection

Create a new connection between two nodes or a node and the flowgraph container.

## Context

A connection represents the directed flow of data from a source to a target. The source and the target of a connection are anchors. The connection defines a table mapping between these table types defined by its source and target. The source of a connection is either an input anchor of the flowgraph container or an output anchor of a node. The target of a connection is either an output anchor of the flowgraph container or an input anchor of a node. An anchor can be the source of several connections. It can only be the target of one connection. A fixed content anchor cannot be the target of any connection.

## Procedure

1. Select without releasing the *Connect* button  in the context button pad of the source anchor of the connection.
2. Drag a connection to the target anchor.

### **i** Note

Depending on the node of the target anchor, the *Create Input Table Mapping* wizard may open. This wizard helps you to choose the right mapping for the connection. You can still change this mapping in the *Mapping Editor* after completing the wizard. To open the wizard again, you have to remove the connection and create it again.

## Results

A new connection between the source anchor and the target anchor is created. If possible, the signature of the source anchor is copied to the target anchor and propagated forward through the flowgraph.

### **i** Note

You can also add a new anchor to an anchor region and create a connection to this anchor in a single action. Instead of dragging the connection to an anchor, drag it to a free position in an anchor region. A new target anchor with the same signature as the source anchor is added before the connection is created.

## Related Information

[Adding an anchor \[page 1307\]](#)

### 9.3.3.4.9 Using the Table Editor

Edit embedded table like anchor signatures and fixed content tables.

## Context

Embedded tables appear in various flowgraph elements. For example, anchors have signature tables and may have fixed content tables. Specialized nodes may have tables in the *General* tab of the *Properties* view. The SAP HANA Application Function Modeler provides a *Table Editor* to edit these tables.

---

## Procedure

- Add, remove, and re-order rows of the embedded table by selecting the respective operations on the right side of the *Table Editor*.
- Edit an entry in the table by double-clicking the respective cell.

## Related Information

[Flowgraphs \[page 1296\]](#)

[Editing an anchor \[page 1308\]](#)

### 9.3.3.4.10 Using the Mapping Editor

Edit the mappings between table types in the *Mappings* tab of the *Properties* view of a flowgraph element.

## Prerequisites

You have selected the *Mappings* tab of the *Properties* view of a flowgraph element.

## Context

A mapping is a projection between table types. The *Mapping Editor* allows you to edit mappings between a number of source and target table types. The left side of the editor shows the source table types, the right side shows the target table types. A binding of two attributes is indicated by a line between them.

#### **i** Note

The mapping editor is used to define the mappings of connections and possible projections within nodes (for example, the Filter node, the Join Node, and the Union Node). It is also used to edit this inputs and outputs of the flowgraph container and of nodes which do not define a projection. In this case, no lines are drawn between the attributes.

#### **i** Note

Not all flowgraph elements allow free editing of all their mappings and table types. In this case the functionality of the *Mapping Editor* is restricted to the permitted editing operations.

## Procedure

- (Optional) To remove a table type, select it and press the minus sign on the right side of the *Mapping Editor*.
- To re-order the source or target table types, click on a table type and use the up/down arrows on the right side of the *Mapping Editor*.
- (Optional) To remove an attribute, select it and press the minus sign on the right side of the *Mapping Editor*.
- (Optional) To re-order the source or target attributes, click on an attribute and use the up/down arrows on the right side of the *Mapping Editor*.
- (Optional) To add an attribute from the source type to the target type, drag the source attribute and drop it on the root of the target tree.  
The attribute is appended at the end of the target attribute list. If the *Mapping Editor* defines a mapping, it is connected by a line with the source attribute indicating an attribute binding.
- (Optional, only available if the *Mapping Editor* defines a mapping) To re-assign a source attribute to a target attribute that is already assigned, drag the source attribute to the target attribute.  
The old binding is replaced by the new one.

## Related Information

[Flowgraphs \[page 1296\]](#)

[Editing the flowgraph container \[page 1301\]](#)

[Editing a node \[page 1305\]](#)

### 9.3.3.4.11 Using the Expression Editor

Compose expressions for filters, join conditions, and calculated attributes.

## Context

The *Expression Editor* allows you to compose SQL expressions based on table type attributes and functions. It consists of an *Function Palette* on the top, an *Attribute Palette* on the left and a *Text Field* on the right.

#### **i** Note

The expression validation is disabled in the SAP HANA Application Function Modeler.

## Procedure

- Type the expression in the *Text Field*.

### **i** Note

Press CTRL + Space bar for auto-completion.

- Select operators and functions in the *Function Palette* to add them to the *Text Field*.
- Drag attributes from the *Attribute Palette* to the *Text Field*.

## Related Information

[Aggregation \[page 1323\]](#)

[Filter \[page 1332\]](#)

[Join \[page 1336\]](#)

## 9.3.3.4.12 Using the Annotation Editor

Add arbitrary annotations to the flowgraph container or a node.

### Context

The flowgraph container and all nodes have an *Annotation* tab in their *Properties* view. Annotations are nested key-value pairs. The SAP HANA Application Function Modeler (AFM) provides an *Annotation Editor* to edit existing annotations like the `sap.afm.palette` annotation or to add your own annotations.

### **i** Note

When the flowgraph is activated, all annotations are exposed in a table with the name extension `.META` in the flowgraph target schema. This way, they can be consumed at runtime.

### **i** Note

For some nodes, the annotations `sap.afm.displayName` and `sap.afm.description` are visible in the *Annotation Editor*. These annotations are for internal use of the AFM and not supposed to be modified.

## Procedure

- Add, remove, and re-order annotations by selecting the respective operations on the right side of the *Annotation Editor*.
- Edit the Key and the Value of an annotation by double-clicking the respective cell.
- Add nested annotations by first selecting an annotation row and then the *Add Child* operation on the right side of the *Annotation Editor*.

---

A nested annotation appears below the selected annotation.

- Collapse and expand nested annotations by selecting the triangle to the left of an annotation key.

## Related Information

[Flowgraphs \[page 1296\]](#)

[Customizing the Node Palette \[page 1319\]](#)

## 9.3.3.5 Tutorial

### Prerequisites

- You have access to a running SAP HANA development system.
- You have a valid user account in the SAP HANA database on that system.
- Your user has been granted the MODELING role.
- Your user has been granted the EXECUTE privilege for the object SYS.REPOSITORY\_REST.
- Your user has been granted the following repository package privileges:
  - repo.read
  - repo.activate\_native\_objects
  - repo.edit\_native\_objects
  - repo.maintain\_native\_packages
- The system user \_SYS\_REPO has SELECT and ALTER privileges on the schema of your user.
- You have access to SAP HANA Studio and opened the *SAP HANA Development* perspective.
- You have created a system in the *System* view in the and logged on to this system with your user.
- You have created a repository workspace for the system.
- You have created a project in the *Project Explorer* view and shared it with the system via the workspace.

#### → Tip

To share a project, right-click on the project and choose ► *Team* ► *Share* ► *Project* ▾ in the context-sensitive menu. In the *Share Project* wizard, choose *SAP HANA Repository* on the first page and choose your repository workspace on the second page.

## Context

This tutorial leads you through the most common steps of using the SAP HANA Application Function Modeler (AFM). At the end of this tutorial, you will have created and tested a runtime procedure with the AFM.

## Procedure

1. Open the SQL console of the system and create the table type WEATHER and the two tables NORTH and SOUTH in your user's schema by executing the following script.

```
CREATE TYPE "WEATHER" AS TABLE ("REGION" VARCHAR(50), "SEASON" VARCHAR(50),
"TEMPERATURE" INTEGER);
CREATE COLUMN TABLE "NORTH" LIKE "WEATHER";
INSERT INTO "NORTH" VALUES ('North', 'Spring', 10);
INSERT INTO "NORTH" VALUES ('North', 'Summer', 23);
INSERT INTO "NORTH" VALUES ('North', 'Autumn', 12);
INSERT INTO "NORTH" VALUES ('North', 'Winter', 2);
CREATE COLUMN TABLE "SOUTH" LIKE "WEATHER";
INSERT INTO "SOUTH" VALUES ('South', 'Spring', 18);
INSERT INTO "SOUTH" VALUES ('South', 'Summer', 34);
INSERT INTO "SOUTH" VALUES ('South', 'Autumn', 23);
INSERT INTO "SOUTH" VALUES ('South', 'Winter', 12);
```

After refreshing the catalog, the table type WEATHER with the three attributes REGION, SEASON, and TEMPERATURE appears in the directory **Procedures > Table Types** of your user's schema. The two tables NORTH and SOUTH with the same signature appear in the directory **Tables** your user's schema.

### NORTH

REGION	SEASON	TEMPERATURE
North	Spring	10
North	Summer	23
North	Autumn	12
North	Winter	2

### SOUTH

REGION	SEASON	TEMPERATURE
South	Spring	18
South	Summer	34
South	Autumn	23
South	Winter	12

2. In the *Project Explorer* view, right-click on the existing project and choose **New > Other** in the context-sensitive menu.  
The *New* wizard appears.
3. Choose **SAP HANA > Database Development > Flowgraph Model**, and then click *Next*.  
The *New Flowgraph Model* wizard appears.
4. In the text field *File Name* enter `avg_temp` as name of the new flowgraph and select *Finish*.  
The system automatically adds the file extension `.hdbflowgraph`. The AFM opens and in the *Editing Area* the empty flowgraph container is displayed.
5. Select the flowgraph container and enter the schema of your user to the *Target Schema* field in the *Properties* view.
6. Add the table NORTH from the *Node Palette* to the flowgraph. For this, drag the Data Source entry from the *General* tab of the *Node Palette* (located on the right side of the AFM) to the flowgraph (choose any free

space inside the canvas of the flowgraph container). Choose the table NORTH from the schema of your user in the dialog that appears.

The node NORTH is added to the flowgraph.

7. Add the table SOUTH from the catalog to the flowgraph. For this, navigate in the catalog to the directory *Tables* in your schema (either in the *Project Explorer* view or in the *Systems* view). Drag the table SOUTH from the catalog to the flowgraph (place it below the NORTH node). Choose *Data Source* in the dialog that appears.  
The node SOUTH is added to the flowgraph.
8. Add a Union node to the flowgraph. For this, drag the Union entry from the *General* tab of the *Node Palette* to the flowgraph (place it right of the other two nodes).  
The node UNION is added to the flowgraph.
9. Create a connection between the DATA anchor of the NORTH node and the INPUT1 anchor of the UNION node. Click the *Connect* button  in the context button pad of the DATA anchor and drag a connection to the INPUT1 anchor.  
A connection between the NORTH node and the UNION node is created.
10. Create a second connection between the DATA\_2 anchor of the SOUTH node and the INPUT2 anchor of the UNION node.  
A connection between the SOUTH node and the UNION node is created.
11. Create a connection between the OUTPUT anchor of the UNION node and the output anchor region of the flowgraph container (the light-blue area at its right boundary).  
The output anchor OUTPUT\_2 is added to the output anchor region of the flowgraph container and a connection between the UNION node and the new anchor is created.
12. Save the flowgraph. Select **File > Save** in the HANA Studio main menu.
13. Activate the flowgraph. For this, right-click the flowgraph object in the *Project Explorer* view and choose **Team > Activate** from the context-sensitive menu.  
A new procedure is generated in the schema of your user.

### Caution

If the system user `_SYS_REPO` does not have SELECT and ALTER privileges then the activation fails.

14. Execute the generated procedure. For this, select the *Execute* button  in the top right corner of the AFM. The *Data Preview* view opens. It contains a tab with the SQL command that calls the generated procedure (with no input and one output) and a tab with the result of the procedure. This result is the union of the tables NORTH and SOUTH.
15. Return to the *AFM* view for the `avg_temp` flowgraph.
16. Add an Aggregation node from the *General* compartment of the *Node Palette* to the flowgraph (place it right of the UNION node).
17. The node AGGREGATION is added to the flowgraph.
18. Connect the OUTPUT anchor of the UNION node with the INPUT anchor of the AGGREGATION node.  
The *Mapping Editor* for the connection is shown in the *Properties* view.
19. In the *Target* area of the *Mapping Editor* for the new connection, select the attribute SEASON of the target INPUT. Remove this attribute by clicking the *Remove* button on the right side of the *Mapping Editor*.  
The attribute SEASON and the corresponding mapping are deleted.
20. Select the AGGRAGATION node. In the *General* tab of its *Properties* view double-click the action of the attribute TEMPERATURE and change it to the value AVG.

21. Create a connection between the OUTPUT\_3 anchor of the AGGREGATION node and the output anchor region of the flowgraph container.  
The output anchor OUTPUT\_4 is added to the output anchor region of the flowgraph container and a connection between the AGGREGATION node and the new anchor is created.
22. Save and activate the flowgraph. Execute the generated procedure.  
The *Data Preview* view opens again. It contains a tab with the SQL command that calls the generated procedure (with no input and two outputs) and two tabs with the results of the procedure. One result is still the union of the tables NORTH and SOUTH. The other result shows in two rows the average temperatures for the regions North and South.
23. Return to the *AFM* view for the avg\_temp flowgraph.
24. Delete the OUTPUT\_2 anchor of the flowgraph container by choosing *Delete* in its context menu (or the respective button in the context button pad).
25. Save and activate the flowgraph. Execute the generated procedure.  
The *Data Preview* view opens again. It contains a tab with the SQL command that calls the generated procedure (with no input and one output) and a second result tab that again shows in two rows the average temperatures for the regions North and South.
26. Return to the *AFM* view for the avg\_temp flowgraph.
27. Delete the SOUTH node from the flowgraph.  
The SOUTH node and its connection to the UNION node is deleted.
28. Create an additional input anchor for the flowgraph by adding the table type WEATHER from the catalog.  
For this, navigate to the directory ► *Procedures* ► *Table Types* ► in the catalog and drag the entry WEATHER to the input anchor region of the flowgraph container.  
The input anchor DATA\_2 is added to the flowgraph.
29. Create a connection between the new DATA\_2 anchor and the INPUT\_2 anchor of the UNION node.  
A new connection between the DATA\_2 anchor and the UNION node is created.
30. Save and activate the flowgraph. Execute the generated procedure.  
A dialog appears where you can choose the free input DATA\_2. Enter the table SOUTH in your user's schema to the Catalog Object field. The *Data Preview* view opens. Again, it contains a tab with the SQL command that calls the generated procedure (with one input and one output) and a second result tab that shows in two rows the average temperatures for the regions North and South.
31. Close the flowgraph. Select ► *File* ► *Close* ► in the HANA Studio main menu.

## Results

You have created a stored procedure that has one input table of the table type WEATHER and one output table that is produced by first forming the union of the table NORTH with the input table and then calculating the average temperature of each season. This procedure can now be used in any application that consumes stored procedures.

## Related Information

[Modeling a flowgraph \[page 1298\]](#)

## 9.3.3.6 Node palette flowgraphs

A node palette flowgraph represents a node palette or a node palette compartment.

The *Node Palette* of the SAP HANA application function modeler is customizable. A custom node palette is represented by a node palette flowgraph. These flowgraphs have the file extension `.hdbflowgraphtemplate` in the *Project Explorer* view.

A node palette flowgraph contains Palette Container and template nodes. These represent the compartments or sub-compartments and the node templates of the corresponding node palette. The Palette Container nodes and template nodes have a nested structure. This structure represents the hierarchy of the corresponding node palette. Moreover, all nodes in a node palette flowgraph are aligned on a horizontal line. Their order (from left to right) represents the order of the node palette entries (from top to bottom).

The *Node Palette* hierarchy can have up to three levels.

1. The first level contains the compartments (for example, the *General* compartment of the application function modeler *Node Palette*). Nodes are not permitted on this level.
2. The second level contains nodes (for example, the Filter node) and sub-compartments.
3. The third level contains only nodes.

A node palette flowgraph represents either a complete node palette or a compartment of the node palette. In the first case, the nesting depth of the node palette flowgraph is at least two and at most three, in the second case, the nesting depth is at most two.

Each flowgraph can be assigned its own custom node palette. This is specified either on creation of the flowgraph or in the *Annotations* tab of the *Properties* view of the flowgraph container.

### Related Information

[Flowgraphs \[page 1296\]](#)

[Editing the flowgraph container \[page 1301\]](#)

### 9.3.3.6.1 Exporting the Node Palette

Export the *Node Palette* as a node template flowgraph.

#### Procedure

1. Right-click the *Node Palette* and choose *Export entire palette* from the context-sensitive menu. The *Save As* wizard appears.
2. Navigate to the directory of your project and save the node template flowgraph file with the extension `.hdbflowgraphtemplate` in this project. Refresh the *Project Explorer* view, and then the node template flowgraph is available in your project.

### **i** Note

The standard location of HANA projects on your local system is the directory `hana_work` in the home directory of your local user. There you find a sub-directory corresponding to the system shared with your project. The directory of the project is then located in the sub-directory `__empty__`.

## 9.3.3.6.2 Customizing the Node Palette

Customize the node palette of a flowgraph by adding a reference to a node palette flowgraph to the annotations of its flowgraph container.

### Context

A flowgraph can be assigned a custom node palette. This can be done in three ways.

- Add additional compartments to the existing AFM node palette.
- Add additional compartments to an empty node palette.
- Add additional compartments to a custom node palette.

### **i** Note

The recommended way to customize the node palette of a flowgraph is via the *New Flowgraph Wizard* during the creation of the flowgraph. The following procedure of directly editing the annotations of the flowgraph container is only advised if you actually need to change the node palette of an existing flowgraph.

### Procedure

1. Open the *Annotations* tab in the *Properties* view of the flowgraph container of a flowgraph.
2. If the annotation does not exist, add the annotation with the key `sap.afm.palette`.
3. (Optional) Insert the name of a node palette flowgraph (with the extension `.hdbflowgraphtemplate`) to the *Value* of this annotation. This replaces the default AFM node palette with the custom node palette defined by the specified node palette flowgraph.
4. If the nested annotation with the key `isDefaultUsed` does not exist, add it as a child to the annotation `sap.afm.palette`.  
The *Value* of this annotation determined if the default AFM node palette is shown.
5. If the nested annotation with the key `additions` does not exist, add it as a child to the annotation `sap.afm.palette`.
6. (Optional) Insert a comma-separated list of names of node palette flowgraphs (with the extension `.hdbflowgraphtemplate`) to the *Value* of this annotation. This adds the compartments defined by the specified node palette flowgraph to the node palette of the flowgraph.

## 9.3.3.6.3 Editing a node palette flowgraph

Edit a node palette flowgraph to model a custom node palette.

### Prerequisites

You have exported the *Node Palette* of the SAP HANA application function modeler to a node palette flowgraph `Template.hdbflowgraphtemplate`.

In addition, you have created a new (standard) flowgraph `Custom.hdbflowgraph` with the advanced option of choosing the node palette flowgraph `Template.hdbflowgraphtemplate` as the Custom Node Palette.

### Context

A node palette flowgraph represents a custom node palette for the application function modeler. Node palette flowgraphs can be edited with the application function modeler like standard flowgraphs. The behavior of the application function modeler when editing node palette flowgraphs differs in two aspects from the editing of standard flowgraphs.

1. All nodes in the node palette flowgraph are automatically aligned on a horizontal line. By this, the order of the nodes (left to right) represents the order of the custom node palette entries (top to bottom).
2. The node palette flowgraph contains nested Palette Container nodes. These nodes represent the hierarchical structure of the custom node palette. These nodes look and behave similar to the flowgraph container.

In the following step by step tutorial, we use the application function modeler to customize the node palette of the Custom flowgraph by editing the Template node palette flowgraph. We cover only those aspects of modeling node palette flowgraphs that differ from modeling standard flowgraphs.

### Procedure

1. Open the Custom flowgraph with the application function modeler.  
The application function modeler displays the empty Custom flowgraph with a custom node palette defined by the Template node palette flowgraph. At this point, this is still the default application function modeler node palette.
2. Open the Template node palette flowgraph with the application function modeler.  
The Template node palette flowgraph is displayed in a separate tab. The flowgraph container contains Palette Container nodes representing the top compartments of the node palette for the Custom flowgraph.

#### **i** Note

A node palette flowgraph contains no connections. Therefore the flowgraph container has no anchor regions. Creating connections is disabled when editing node palette flowgraphs.

3. Right-click the GENERAL node and choose *Collapse/Expand* in the context-sensitive menu. The GENERAL node expands. It contains the template nodes of the General compartment of the application function modeler node palette.

#### **i** Note

You can collapse a Palette Container node by choosing again *Collapse/Expand* in the context-sensitive menu.

4. Drag the JOIN node to a position between the SORT node and the UNION node. The auto-layout function of the application function modeler rearranges the nodes such that the JOIN node and the SORT node have effectively swapped positions.
5. Switch to the editing tab of the Custom flowgraph. Refresh the custom *Node Palette* by right-clicking the *Node Palette* and choosing *Refresh* in the context-sensitive menu. The Join node template and the Sort node template have swapped places in the General compartment of the *Node Palette*.
6. Switch to the editing tab of the Template node palette flowgraph. Add a Palette Container node to the GENERAL node by dragging the corresponding node template from the *General* compartment of the *Node Palette* to the canvas of the GENERAL node. A nested Palette Container node named COMPARTMENT is added to the GENERAL node.
7. Add an object from the *Project Explorer* view to the canvas of the COMPARTMENT node.
8. Switch to the editing tab of the Custom flowgraph. Refresh the custom *Node Palette*. The sub-compartment *Palette Container* is added to the *General* compartment of the custom *Node Palette*. It contains the node template for the object from the *Project Explorer* view added to the COMPARTMENT node in the previous step.
9. Switch to the editing tab of the Template node palette flowgraph. Add a Filter node from the *Node Palette* to the COMPARTMENT node. Edit the Display Name and the Description in the *General* tab of the *Properties* view of the Filter node. In addition, edit the signatures of the input and the output of the Filter node and define a filter expression.
10. Switch to the editing tab of the Custom flowgraph. Refresh the custom *Node Palette*. A new node template with the chosen display name and description (tool-tip) was added to the Palette Container sub-compartment.
11. Add node template of the new filter node from the custom *Node Palette* to the Custom flowgraph. The added Filter node has received the modified input and output signatures and the filter expression of the Filter node in the Template node palette flowgraph.
12. Switch to the editing tab of the Template node palette flowgraph. Move the COMPARTMENT node from the canvas of the GENERAL node to the canvas of the flowgraph container.
13. Switch to the editing tab of the Custom flowgraph. Refresh the custom *Node Palette*. The previous Palette Container sub-compartment in the General compartment is now a new top level compartment of the *Node Palette*.

## Related Information

[Modeling a flowgraph \[page 1298\]](#)

## 9.3.4 Node Reference

Use the nodes in the General palette to transform the data.

Select the nodes that you want to use and place them on the canvas. Double-click the nodes to begin configuring them. Attach the nodes to other nodes by dragging the arrow from the previous node to the next node.

### 9.3.4.1 AFL Function

Use application functions to perform data intensive and complex operations on a database.

#### Prerequisites

You have existing functions in the Application Function Library (AFL).

#### Context

Typically, businesses create a library of application functions that they use on their databases. Application functions are like database procedures written in C++ and called from outside to perform data intensive and complex operations. These functions are processed in the database, rather than at the application level. You can call one or more of these functions and use them as a part of your flowgraph. Use this node to model functions of the Application Function Library that are registered with the system.

AFL functions are grouped by function areas. Those function areas only support certain kinds of data.

Function Area	Kind
AFLPAL (Predictive Analytics Library)	TABLE
AFLBFL (Business Function Library)	TABLE, SCALAR, and COLUMN

**i Note**  
The COLUMN kind is a table with one column.

You can retrieve the list of all AFL areas and functions registered in a HANA system by viewing the content of the views "SYS"."AFL\_AREAS" and "SYS"."AFL\_FUNCTIONS".

#### **i Note**

The AFL Function node is not available for real-time processing.

## Procedure

1. Select the AFL Function node and place it on the canvas. The *Select a Function* dialog appears.
2. In the *Area* option, select the group that contains the function.
3. Select the *Function*, and then click *OK*.
4. Connect the previous node and select the input port that you want to use, and then click *OK*.
5. Double-click the node to configure the options.
6. To define the argument, select the input port, and then the *Pencil* icon under *Function Parameters*. The *Select a Column* dialog opens.
7. Select one or more columns, and then click *OK*. The columns are added under the *Argument* column.
8. (Optional) If you want to create fixed content columns rather than using columns from an upstream node or data source, do not connect any input source to this port, and complete these substeps.
  - a. In the *Attributes* tab, click the **+** icon to add a column.
  - b. In the *Name* option, enter a unique column name, and choose a *Data Type* from the list. Depending on the data type, you may need to enter a *Length* value also.
  - c. Select the *Nullable* checkbox if the column can have empty values. Click *OK*.
  - d. Repeat these steps to create additional columns. Click the *Fixed Content* tab and then click the *Fixed Content* checkbox.
  - e. Click the **+** icon to add values to each column, and then click *OK*. Repeat this step to add more rows of data.
9. Click *Save* and *Back* to return to the flowgraph editor.
10. Connect the output pipe to the next node or the Data Sink node. The *Select Output Parameter* dialog appears.
11. In the *Select existing* option, choose an output port, and then click *OK*.

## Related Information

[Using the Table Editor \[page 1310\]](#)

### 9.3.4.2 Aggregation

An *Aggregation* node represents a relational group-by and aggregation operation.

## Prerequisites

You have added an Aggregation node to the flowgraph.

#### **i** Note

The Aggregation node is available for realtime processing.

## Procedure

1. Select the Aggregation node.
2. Map the input columns and output columns by dragging them to the output pane. You can add, delete, rename, and reorder the output columns, as needed. To multi-select and delete multiple columns use CTRL/Shift keys, and then click [Delete](#).
3. In the [Aggregations](#) tab, specify the columns that you want to have the aggregate or group-by actions taken upon. Drag the input fields and then select the action from the drop-down list.
4. (Optional) Select the [Having](#) tab to run a filter on an aggregation function. Enter the expression. To view the options in the expression editor, click [Load Elements & Functions](#). You can drag and drop the input and output columns from the [Elements](#) pane, then drag an aggregation function from the [Functions](#) pane. Click or type the appropriate operators. For example, if you want to find the transactions that are over \$75,000 based on the average sales in the 1st quarter, your expression might look like this:

```
AVG("Aggregation1_Input"."SALES") > 75000.
```

Option	Description
<b>Avg</b>	Calculates the average of a given set of column values.
<b>Count</b>	Returns the number of values in a table column.
<b>Group-by</b>	Use for specifying a list of columns for which you want to combine output. For example, you might want to group sales orders by date to find the total sales ordered on a particular date.
<b>Max</b>	Returns the maximum value from a list.
<b>Min</b>	Returns the minimum value from a list.
<b>Sum</b>	Calculates the sum of a given set of values.

5. (Optional) Select the [Filter Node](#) tab to compare the column name against a constant value. Click [Load Elements & Functions](#) to populate the Expression Editor. Enter the expression by dragging the column names, the function, and entering the operators from the pane at the bottom of the node. For example, if you want to the number of sales that are greater than 10000, your expression might look like this:  
"Aggregation1\_input"."SALES" > 10000. See the "SQL Functions" topic in the *SAP HANA SQL and System Views Reference* for more information about each function.
6. Click [Save](#) to return to the Flowgraph Editor.

## Related Information

[Using the Mapping Editor \[page 1311\]](#)

## 9.3.4.2.1 Aggregation Options

Description of options for the Aggregation node.

Option	Description
Name	The name of the node.
Display Name	<p><b>i Note</b> AFM only.</p> <p>The name shown in the Palette pane.</p> <p><b>i Note</b> This option can only be changed when creating a template. It cannot be changed when using the node outside of a template.</p>
Description	<p><b>i Note</b> AFM only.</p> <p>(Optional.) Provides a comment about the operation. For example, "Calculate total sales in May."</p>
Column/Attribute	The input column name that you want to use in an Aggregation operation.
Aggregation/Action	<p>Choose one of the following:</p> <p>Avg: calculates the average of a given set of column values.</p> <p>Count: returns the number of values in a table column.</p> <p>Group-by: use for specifying a list of columns for which you want to combine output. For example, you might want to group sales orders by date to find the total sales ordered on a particular date.</p> <p>Max: returns the maximum value from a list.</p> <p>Min: returns the minimum value from a list.</p> <p>Sum: calculates the sum of a given set of values.</p>

## 9.3.4.3 Data Sink

Edit nodes that represent data sinks.

### Procedure

1. Drag the Data Sink node onto the canvas.
2. In the *Select an Object* dialog, type the name of the object to add, or browse the object tree to select one or more objects, and click *OK*.
3. (Optional) You can click the magnifying glass icon to preview the existing data (if any) in the table. The data will change after the flowgraph runs.
4. In the *General* tab of the *Properties* view use the drop-down menus *Authoring Schema* and *Catalog Object* to specify the data sink.

#### → Tip

You can configure the authoring schema by choosing *Schema Mapping* in the *Quick* view of the *SAP HANA Modeler* perspective.

5. Select *Truncate Table* to clear the table before inserting data. Otherwise, all inserted data is appended to the table.
6. Optionally, if the node is a Data Sink (Template Table) node, specify in the same tab in the drop-down menu *Data Layout* whether a table with row or column layout is created.
7. To optionally create a separate target table that tracks the history of changes, set the *History Table Settings* options.

### Results

The signature of the input anchor is set automatically.

#### i Note

To activate the flowgraph, the database user `_SYS_REPO` needs `INSERT` and in case of truncation also `DELETE` object privileges for the chosen data sink.

### Related Information

[Setting up the SAP HANA Application Function Modeler \[page 1294\]](#)

## 9.3.4.3.1 Data Sink Options

Description of options for the Data Sink node.

Option	Description
Enter table or view name	<p><b>i Note</b> AFM only.</p> <p>Enter the name of the table or view.</p>
Matching items	<p><b>i Note</b> AFM only.</p> <p>Shows matching tables or views as you begin typing in the previous option.</p>

Option	Description
Name	The name for the output target.
Display Name	<p><b>i Note</b> AFM only.</p> <p>The name shown in the Palette pane.</p> <p><b>i Note</b> This option can only be changed when creating a template. It cannot be changed when using the node outside of a template.</p>
Description	<p><b>i Note</b> AFM only.</p> <p>(Optional.) Provides a comment about the target. For example, "West Region Sales Q1."</p>
Type	Lists whether it is a view or table.
Authoring Schema	Lists the system or folder where the view or table is located.
Catalog Object	Lists the table or view.
Truncate Behavior	<p>Limits the amount of data written to the Data Sink.</p> <p>In the SAP HANA Web-based Development Workbench, for the <a href="#">Truncate Table</a> option, select it to clear the table before inserting data. Otherwise, all inserted data is appended to the table.</p>

Option	Description
Writer Type	Choose from the following options: insert: adds new records to a table. upsert: if a record doesn't currently exist, it is inserted into a table. If the record exists, then it is updated. update: includes additional or more current information in an existing record.
Key Generation Attribute	Generates new keys for target data starting from a value based on existing keys in the column you specify.
Sequence Schema	When generating keys, select the schema where the externally created sequence file is located.
Sequence Name	When generating keys, select the externally created sequence to generate the new key values.
Change time column name	Select the target column that will be set to the time that the row was committed. The data type must be TIMESTAMP.
Change type column name	Select the target column that will be set to the row change type. The data type is VARCHAR(1).

History Table Settings	Description
Schema Name	(Optional) Select the schema location where you want the history table.
Name	Enter a name for the history table.
Type	Select whether the target is a database table or a template table.
Data Layout	Specify whether to load the output data in columns or rows.

## Related Information

[Load Behavior Options for Targets in Flowgraphs \[page 1288\]](#)

### 9.3.4.4 Data Source

Edit nodes that represent data sources.

## Prerequisites

You added a Data Source node to the flowgraph.

---

To activate the flowgraph, the database user `_SYS_REPO` needs SELECT object privileges for the chosen data source.

## Procedure

1. Drag the Data Source node onto the canvas.

You can click the magnifying glass icon to preview the existing data in the table or view.

2. In the *Select an Object* dialog, type the name of the object to add, or browse the object tree to select one or more objects, and click *OK*.
3. In the *General* tab of the *Properties* view, use the drop-down menus *Authoring Schema* and *Catalog Object* to specify the data source.

### Note

The check-box *Realtime Behavior* is only relevant if the flowgraph uses the additional cost SAP HANA smart data integration and SAP HANA smart data quality optional component and if a task plan is generated.

### Tip

You can configure the authoring schema by choosing *Schema Mapping* in the *Quick* view of the *SAP HANA Modeler* perspective.

## Results

The signature of the output anchor is set automatically.

## Related Information

[Setting up the SAP HANA Application Function Modeler \[page 1294\]](#)

[Add a Variable to the Flowgraph \[page 1280\]](#)

## 9.3.4.4.1 Data Source Options

Description of the options in the Data Source node.

Option	Description
Name	The name for the node.
Display Name	<p><b>i Note</b> AFM only.</p> <p>The name shown in the Palette pane.</p> <p><b>i Note</b> This option can only be changed when creating a template. It cannot be changed when using the node outside of a template.</p>
Description	<p><b>i Note</b> AFM only.</p> <p>(Optional.) Provides a comment about the source. For example, "West Region Sales Q1."</p>
Type	Lists whether the data source is a view or table.
Authoring Schema	Lists the system or folder where the view or table is located.
Catalog Object	Lists the repository where the table or view is located
Realtime Behavior	Select to run in batch or real-time mode.
Partition Type	<p><b>i Note</b> Web-based Development Workbench only.</p> <p>Choose one of the following:</p> <p>None: does not partition the table</p> <p>Range: divides the table data into sets based on a range of data in a row.</p> <p>List: divides the table into sets based on a list of values in a row.</p>
Attribute	<p><b>i Note</b> Web-based Development Workbench only.</p> <p>The column name used for the partition.</p>

Option	Description
Partition name	<p><b>i Note</b></p> <p>Web-based Development Workbench only.</p> <p>The name for the partition such as "region".</p>
Value	<p><b>i Note</b></p> <p>Web-based Development Workbench only.</p> <p>The range or list.</p>

### 9.3.4.4.2 Reconcile Changes in the Source

The Data Source node in a flowgraph lets you reconcile differences that have occurred between the source table and its current structure in the node.

#### Context

When the structure of the underlying object (SAP HANA table, virtual table, and so on) for a Data Source changes, you can view and update what has changed in the structure since adding the Data Source to the flowgraph.

#### Procedure

1. View the flowgraph containing the Data Source node.
2. Select the *Compare Tables* icon next to the data source.  
The Compare Tables window displays the difference(s) between the old and new versions.
3. To update the flowgraph, select *Reconcile*.
4. Save the flowgraph.

## 9.3.4.5 Filter

A *Filter* node represents a relational selection combined with a projection operation. It also allows calculated attributes to be added to the output.

### Prerequisites

You have added a Filter node to the flowgraph.

#### **i** Note

The Filter node is available for real-time processing.

### Context

#### Web-based Development Workbench

1. Drag the Filter node onto the canvas, and connect the source data or the previous node to the Filter node.
2. Double-click the Filter node.
3. (Optional) Enter a name for this Filter node in the *Node Name* option.
4. (Optional) Select *Distinct* to output only unique records.
5. (Optional) To copy any columns that are not already mapped to the output target, drag them from the Input pane to the Output pane. You may also remove any output columns by clicking the pencil icon or the trash icon, respectively. You can multi-select the columns that you do not want output by using the CTRL or Shift key, and then Delete.
6. (Optional) Click *Load Elements & Functions* to populate the Expression Editor. Drag input columns into the *Mapping* tab to define the output mapping and perform some sort of calculation. Choose the functions and the operators. For example, you might want to calculate the workdays in a quarter, so you would use the *Workdays\_Between* function in an expression like this: `WORKDAYS_BETWEEN (<factory_calendar_id>, <start_date>, <end_date> [, <source_schema>])`. Click *Validate Syntax* to ensure that the expression is valid.
7. Click the *Filter node* tab and then click *Load Elements & Functions* to populate the Expression Editor. You can use the Expression Editor or type an expression to filter the data from the input to the output. Drag the input columns, select a function and the operators. For example, if you want to move all the records that are in Canada, your filter might look like this: `"Filter1_input"."COUNTRY" = "Canada"`. See the "SQL Functions" topic in the *SAP HANA SQL and System Views Reference* for more information about each function.
8. Click *Save* to return to the flowgraph.

#### Application Function Modeler

1. Select the Filter node.
2. Select the *General* tab of the *Properties* view.
3. Select the *Value Help* and use the *Expression Editor* to configure the *Filter Expression*.

4. Add additional attributes for calculated outputs in the *Output* tab.
5. Select the *Mappings* tab. In the *Mapping Editor*, define the output mapping of the node. In addition you can define the calculated attributes by first selecting the attribute in the *Target* list and then selecting *Edit Expression*.  
The *Expression Editor* opens to edit the expression that calculates the attribute.

**i Note**

You need to manually set the type of the calculated attribute.

6.

**Example**

Let's say that you have a single input source, and connected it to a Match node. You selected Most Recent as your survivor rule, so that the output from Match has a Group\_Master column. Those duplicate records with the most recent Last\_Updated date are marked with a value of "M". After connecting the Match node to the Filter node, you can use the following expression to output only the master and unique records:

**Sample Code**

```
("Filter1_Input"."GROUP_ID" is null) OR ("Filter1_Input"."GROUP_ID" is not null and "Filter1_Input"."GROUP_MASTER" = 'M')
```

Prior to the Filter node, some example data might look like the following.

Data input to the Filter node

GROUP_ID	RE-VIEW_GROUP	CON-FLICT_GROUP	LAST_UP-DATED	ADDRESS	ADDRESS2	GROUP_MAS-TER
<null>	<null>	<null>	<null>	1411 Broadway	New York 10018	<null>
<null>	<null>	<null>	<null>	3 Fleetwood Dr	Newberg NY 12550	<null>
<null>	<null>	<null>	<null>	300 Cliffside Dr	Atlanta GA 30350	<null>
1	N	C	01/01/16	332 Front St	La Crosse WI 54601	M
1	N	C	03/10/11	332 Front St	La Crosse WI 54601	<null>
1	N	C	07/04/15	332 Front St	La Crosse WI 54601	<null>
<null>	<null>	<null>	<null>	3738 North Fraser Way	Burnaby BC V3N 1E4	<null>

After the Filter node, you can see that two duplicate entries were removed, and only the master record and the other four unique records are output.

Data output from the Filter node

GROUP_ID	RE-VIEW_GROUP	CON-FLICT_GROUP	LAST_UP-DATED	ADDRESS	LASTLINE	GROUP_MAS-TER
<null>	<null>	<null>	<null>	1411 broadway	new york 10018	<null>
<null>	<null>	<null>	<null>	3 Fleetwood Dr	Newberg NY 12550	<null>
<null>	<null>	<null>	<null>	300 the cliffsup	atlanta 30350	<null>
1	N	C	01/01/16	332 Front st	La Crosse 54601	M
<null>	<null>	<null>	<null>	3738 NORTH FRASER WAY TH 6203	BURNABY BC	<null>

## Related Information

[Using the Mapping Editor \[page 1311\]](#)

[Using the Expression Editor \[page 1312\]](#)

### 9.3.4.5.1 Filter Options

Description of options for the Filter node.

Option	Description
Name	The name for the node.
Display Name	<p><b>i Note</b> AFM only.</p> <p>The name shown in the Palette pane.</p> <p><b>i Note</b> This option can only be changed when creating a template. It cannot be changed when using the node outside of a template.</p>

Option	Description
Description	<p><b>i Note</b> AFM only.</p> <p>(Optional.) Provides a comment about the node. For example, "Only European Data."</p>
Distinct	<p><b>i Note</b> Web-based Development Workbench only.</p> <p>(Optional). Select to output only unique records. The records must match exactly. If you know that you have duplicates, but have a ROW_ID column, or another column that has a unique identifier for each record, then you will want to suppress that column in the Filter node.</p>
Filter Node	<p>Enter an expression so that only the valid records are output based on the expression criteria. You can enter some SQL statements to set the value of the target column. Any of the SAP HANA SQL functions can be used. See the <i>SAP Hana SQL and System Views Reference</i>.</p> <p><b>i Note</b> In AFM, you can use the Expression Editor to assist in creating the expression.</p>

### 9.3.4.6 Input Type

Input Type is used to set parameters for use in the data source tables when the flowgraph is activated.

Before using the Input Type node, you must have an existing table or table type created. You can create an input table type in application function modeler in SAP HANA studio.

You can use the Input Type node to specify the physical table at run time and make it more flexible by setting parameters. The schema of the input tables must match the schema of Input Type.

When you drag the Input Type node into the *Input Types* pane of a flowgraph, in the *Select an Object* dialog, type the name of the object to add, or browse the object tree to select one or more objects, and click *OK*.

#### General Properties

Option	Description
Kind	Table: Cannot be changed.

Option	Description
Real-time behavior	<p>Choose to run as real-time or batch processing. When selecting to run as a real time process, then you must select a Reference Virtual Table.</p> <div style="background-color: #fff9c4; padding: 5px;"> <p><b>i Note</b></p> <p>When choosing real-time processing, you must use the Output Type node. When using batch processing, you can use the Output Type node, the Data Sink node, or the Template Table node.</p> </div>
Reference Virtual Table	Browse to the schema and select the table name. Only used when Real-time behavior is selected.

When executing the flowgraph, you are prompted with the Table Type Parameters window to specify the physical table name. The schema of the input table type and the physical table must match.

#### Example

Let's say that you have many employee tables. The tables are listed by department and all of the schemas are the same with the same columns for employee names, addresses, ID numbers, and so on. You want to replicate the tables from one system and place them into another, and also to cleanse the data in the process. You can use Input Type to pull the tables into one flowgraph by calling the individual tables at run time, whereas if you used the Data Source node, you would have to run a separate flowgraph for each department.

## Related Information

[Output Type \[page 1338\]](#)

### 9.3.4.7 Join

A Join node represents a relational multi-way join operation.

## Prerequisites

You have added a Join node to the flowgraph.

#### **i Note**

The Join node is not available for real-time processing.

## Context

The Join node can perform multiple step joins on two or more inputs.

## Procedure

1. Select the Join node.
2. (Optional) Add additional input anchors.
3. (Optional) Remove any output columns by clicking the pencil icon or the trash icon, respectively. You can multi-select the columns that you do not want output by using the CTRL or Shift key, and then Delete. The Mapping column shows how the column has been mapped with the input source.
4. In the *Properties* view, select the *General* tab to configure the type of the join (inner join, left outer join, or right outer join).
5. In the table defined in the *General* tab, use the *Table Editor* to define the *Left* join partner, the *Join Type*, the *Right* join partner and the *Join Condition* of each join step. In this, only the first entry in the join condition consists of a *Left* join partner and a *Right* join partner. Every subsequent join condition has the previous join tree as *Left* join partner.  
The *Expression Editor* opens and lets you specify the *Join Condition*.
6. In the *Mappings* tab, use the *Mapping Editor* to edit the output attributes of the join.

## Related Information

[Using the Table Editor \[page 1310\]](#)

[Using the Mapping Editor \[page 1311\]](#)

[Using the Expression Editor \[page 1312\]](#)

[Adding an anchor \[page 1307\]](#)

### 9.3.4.7.1 Join Options

Description of options for the Join node.

Option	Description
Name	The name for the node.

Option	Description
Display Name	<p><b>i Note</b> AFM only.</p> <p>The name shown in the Palette pane.</p> <p><b>i Note</b> This option can only be changed when creating a template. It cannot be changed when using the node outside of a template.</p>
Description	<p><b>i Note</b> AFM only.</p> <p>(Optional.) Provides a comment about the node. For example, "Employee_v8 and Employee_v12."</p>
Left	The left source of a join.
Join Type	<p>Choose from one of these options:</p> <p>Inner: use when each record in the two tables has matching records.</p> <p>Left_Outer: output all records in the left table, even when the join condition does not match any records in the right table.</p> <p>Right_Outer: output all records in the right table, even when the join condition does not match any records in the left table.</p>
Right	The right source of a join.
Join Condition	<p>The expression that specifies the criteria of the join condition.</p> <p><b>i Note</b> In AFM, you can use the Expression Editor to assist in creating the expression.</p>
Add	A join condition is created.
Remove	The highlighted join condition is deleted.

### 9.3.4.8 Output Type

Output Type is used to set parameters for use in the output tables when the flowgraph is activated.

Before using the Output Type node, you must have an existing table or table type created. You can create an output table type in application function modeler within SAP HANA Studio.

**i Note**

If you use an Input Type node and select real-time processing, you must use the Output Type node.

---

When you drag the Output Type node into the *Output Types* pane of a flowgraph, in the *Select an Object* dialog, type the name of the object to add, or browse the object tree to select one or more objects, and click *OK*.

## General Properties

Option	Description
Kind	Table: Cannot be changed.

## Related Information

[Input Type \[page 1335\]](#)

### 9.3.4.9 Procedure

Use the Procedure node when you want to invoke a SAP HANA procedure or a virtual procedure.

## Prerequisites

- To activate the flowgraph, the database user `_SYS_REPO` needs the EXECUTE object privilege on the procedure to be selected.
- If a procedure has scalar parameters, the values for these parameters come from variables defined in the flowgraph. These variables are created automatically. When the flowgraph is executed, you provide a value for each variable so that it can be passed to the input scalar parameters.

## Context

### **i** Note

The Procedure node is not available for real-time processing.

## Procedure

1. Drag the Procedure node onto the canvas.

2. In the *Select an Object* dialog, type the name of the object to add, or browse the object tree to select one or more objects, and click *OK*.
3. Select the Procedure node.
4. The following step applies only if you added the Procedure node from the *Node Palette*.
  - In SAP HANA studio, in the *General* tab of the *Properties* view, select the drop-down menus for the *Schema* and the *Procedure* that is represented by the node.
  - In SAP HANA Web-based Development Workbench, open the node and select a *Schema Name* and the *Procedure Name* for the node.
5. Add one or more inputs to the node.
6. Open the Procedure node and map the input parameters:
  - a. Select the parameter name in the top pane.
  - b. On the *Schema* tab in the bottom pane, map each column to an *Input Mapping Port Column* using the drop-down list for each column.
  - c. Save the flowgraph.

### 9.3.4.9.1 Procedure options

Description of options for the Procedure node.

Option	Description
Name	The name for the node.
Display Name	<p><b>i Note</b> AFM only.</p> <p>The name shown in the Palette pane.</p> <p><b>i Note</b> This option can only be changed when creating a template. It cannot be changed when using the node outside of a template.</p>
Description	<p><b>i Note</b> AFM only.</p> <p>(Optional.) Provides a comment about the node. For example, "Run schedule."</p>
Schema	The location and definition of the procedure.
Procedure	The stored procedure that you want to run in the flowgraph.

## 9.3.4.10 R-Script

Use the R-Script node for developing and analyzing statistical data.

R is an open-source programming language and software environment for statistical computing and graphics. The R code is embedded in SAP HANA SQL code in the form of a RLANG procedure. You can embed R-function definitions and calls within SQL Script and submit the code as part of a query to the database. You can also use R-Script to define the dataflow and schedule flowgraph processing.

### **i** Note

The R-Script node is not available for real-time processing.

Only table types are supported in the R-Script node.

1. Place the R-Script node onto the canvas and connect the previous node. The *Select Input Parameter* dialog appears.
2. Choose one of the options, and then click *OK*.
  - *Select existing*: select the input port name that you want to use. These inputs are predefined in the node.
  - *Create new*: creates an additional input port where you can enter fixed content or connect another input port.
3. Double-click the node to configure the options. The *Script* tab shows a predefined template where you can enter your R-script code.
4. Click the *Parameters* tab. The defined IN/OUT ports and whether they are connected are shown. If you have connected this node to an upstream node, click the connected IN port to see the column information in the *Attributes* table.
5. (Optional) To create additional input or output ports, click *Add* in the *Parameters* table. Enter a unique name and choose whether it is an input or output port. Click *OK*.
6. (Optional) To create fixed content columns rather than using the columns from an upstream node or data source, do not connect any input source to this port, and complete these substeps.
  1. In the *Attributes* table, click *Add*.
  2. In the *Name* option, enter a unique column name, and choose a *Data Type* from the list. Depending on the data type, you may need to enter a *Length* value also.
  3. Select the *Nullable* checkbox if the column can have empty values. Click *OK*.
  4. Repeat these substeps to create additional columns.
  5. Select the *Fixed Content* checkbox.
  6. Click the **+** icon to add values to each column. Repeat this step to add more rows of data.
7. After you have finished configuring the input and output ports, click *Save* and *Back* to return to the flowgraph editor.
8. Connect the output port to the next node. The *Select Output Parameter* dialog appears.
9. In the *Select existing* option, choose an output port, and then click *OK*.

## Related Information

[SAP HANA R Integration Guide \(HTML\)](#)

## 9.3.4.11 Sort

A Sort node represents a relational sort operation.

### Prerequisites

You have added a Sort node to the flowgraph.

### Context

The Sort node performs a sort by one or more attributes of the input.

#### **i** Note

The Sort node is available for real-time processing.

### Procedure

1. Select the Sort node.
2. In the *Properties* View, select the *General* tab to configure the sort order.
3. In the *General* tab, use the *Table Editor* to define the *Attributes* and the *Sort Order* by which the input is sorted. It is possible to specify several *Attributes* with descending priority.

### Related Information

[Using the Table Editor \[page 1310\]](#)

#### 9.3.4.11.1 Sort Options

Description of options for the Sort node.

Option	Description
Name	The name for the node.

Option	Description
Display Name	<p><b>i Note</b> AFM only.</p> <p>The name shown in the Palette pane.</p> <p><b>i Note</b> This option can only be changed when creating a template. It cannot be changed when using the node outside of a template.</p>
Description	<p><b>i Note</b> AFM only.</p> <p>(Optional.) Provides a comment about the node. For example, "Sort ascending sales order."</p>
Column/Attribute	The column used for sorting.
Sort Type/Sort Order	<p>How to sort the data.</p> <p>Ascending: When sorting numerical data, put the smallest number first. When sorting alphabetically, start with the first letter.</p> <p>Descending: When sorting numerical data, put the largest number first. When sorting alphabetically, start with the last letter.</p>
Add	A row is configured to be used for sorting.
Remove	The highlighted entry is deleted, so that it will not be used in sorting.
Up	The entry is moved up so that it is sorted before any entries below it.
Down	The entry is moved down so that it is sorted after any entries above it.

### 9.3.4.12 Union

A Union node represents a relational union operation.

#### Prerequisites

You have created a Union node in the flowgraph.

## Context

The union operator forms the union from two or more inputs with the same signature. This operator can either select all values including duplicates (UNION ALL) or only distinct values (UNION).

### **i** Note

The Union node is available for real-time processing.

## Procedure

1. Select the Union node.
2. (Optional) Add additional input anchors.
3. In the *General* tab of the *Properties* view define whether the operator is a UNION ALL or a UNION operator by selecting or unselecting the checkbox *Create Union All*.

## Related Information

[Adding an anchor \[page 1307\]](#)

### 9.3.4.12.1 Union Options

Description of options for the Union node.

Option	Description
Name	The name for the node.
Display Name	<p><b>i</b> Note</p> <p>AFM only.</p> <p>The name shown in the Palette pane.</p> <p><b>i</b> Note</p> <p>This option can only be changed when creating a template. It cannot be changed when using the node outside of a template.</p>

Option	Description
Description	<p data-bbox="667 365 772 398"><b>i Note</b></p> <p data-bbox="667 416 759 443">AFM only.</p> <p data-bbox="647 483 1350 539">(Optional.) Provides a comment about the node. For example, "Combine HR2015 and HR2010."</p>
Create Union All	<p data-bbox="647 562 1358 618">The option to merge all of the input data (including duplicate entries) into one output, when selected.</p>

# 10 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.

## 10.1 SAP HANA HDBSQL Options

Execute SAP HANA HDBSQL commands to query information about the database and database objects.

### **i** Note

In addition to SAP HANA HDBSQL commands, you can also enter an SQL statement or a database procedure. The statement or procedure must be in quotation marks.

### Configuration Options

Use the following options to modify the operation of SAP HANA HDBSQL commands.

Database Session

Option	Description
<code>-i &lt;instance_number&gt;</code>	Specifies the instance number of the system
<code>-n &lt;host&gt;[:&lt;port&gt;]</code>	Specifies the name of the computer on which the system is installed and the port number
<code>-d &lt;database_name&gt;</code>	Specifies the name of the multitenant database container in a multiple-container system
<code>-u &lt;database_user&gt;</code>	Specifies the user name for logging on to the database
<code>-p &lt;database_user_password&gt;</code>	Specifies the password for logging on to the database
<code>-U &lt;user_store_key&gt;</code>	Uses credentials from the user store
<code>-e</code>	Specifies that encrypted data transmission is used

Option	Description
-r	Enforces the execution of SQL statements as statements rather than as prepared statements
-S <sql mode>	Specifies the SQL mode, either "INTERNAL" or "SAPR3"
-z	Switches off AUTOCOMMIT mode
-r	Suppresses usage of prepared statements
-saml-assertion <file>	Uses a file to provide a SAML assertion

#### Input and Output

Option	Description
-c <separator>	Specifies the separator used to separate individual commands when importing commands from a file. The default value is ;
-I <file>	Imports commands from a batch file
-m	Activates multiple line mode for entering SAP HANA HDBSQL commands
-o <file>	Writes the results to a file
-x	Suppresses additional output such as the number of selected rows in a result set
-resultencoding <encoding>	Forces output encoding for result data: can be one of <b>UTF8</b> , <b>LATIN1</b> or <b>AUTO</b> (the default)

#### Formatting Output

Option	Description
-A	Returns the result set in an aligned format
a	Suppresses the output of the column names in the result set
-C	Suppresses escape output format
-b <maximum_length>	Defines the maximum number of characters for output of LOB values (the default value is 10)
-f	Returns all SQL statements that are sent to the database instance
-F <separator>	Specifies which string SAP HANA HDBSQL uses as a separator between the individual columns of the result set (the default value is  )
-g <>null_value>	Specifies the character for NULL values in the result set (the default value is ?)
-p <prefix>	Specifies which string is to be output before each row of the result set (the default value is  )
-P <suffix>	Specifies which string is to be output after each row of the result set (the default value is  )
-Q	Outputs each column of the result set in a new row

Option	Description
-j	Switches off the page by page scroll output

#### Other

Option	Description
-h	Displays the help
-t	Outputs debug information
-T <file>	Activates the SQLDBC trace, which writes the trace data to the specified file
-v	Displays version information about the SAP HANA HDBSQL program

#### SSL Options

Option	Description
-sslprovider <provider>	Specifies the cryptographic service provider used for SSL connections (one of commoncrypto, sapcrypto, mscrypto)
-sslkeystore	Specifies the SSL keystore name
-ssltruststore	Specifies the SSL truststore name
-ssltrustcert	Skips certificate validation
-sslhostnameincert	Specifies the hostname that is used for certificate validation
-sslcreatecert	Creates a self-signed certificate

## Interactive Options

Use the following options when operating SAP HANA HDBSQL in interactive mode.

Command	Description
\?	Displays all HDBSQL commands
\h[elp]	
\a[utocommit] [ON OFF]	Switches AUTOCOMMIT mode on or off
\a[ign] [ON OFF]	Controls whether SQL statement results are formatted
\e[scape] [ON OFF]	Switches the escape output format on or off
\c[onnect]	Logs a user onto the database

Command	Description
\dc [PATTERN]	<p>Lists all table columns that correspond to the specified [PATTERN] and to which the current user has access.</p> <p>[PATTERN] can be specified as follows: [SCHEMA.] [OBJECT_NAME]. The following placeholders are possible:</p> <ul style="list-style-type: none"> <li>• For one character: _</li> <li>• For any number of characters: %</li> </ul> <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"> <li>• Column name</li> <li>• Data type</li> <li>• Column length</li> <li>• Null value permitted or not</li> <li>• Position of column in primary key of table (if applicable)</li> </ul>
\de [PATTERN]	<p>Lists all the indexes of database objects that correspond to the specified [PATTERN].</p> <p>[PATTERN] can be specified as follows: [SCHEMA.] [OBJECT_NAME]. The following placeholders are possible:</p> <ul style="list-style-type: none"> <li>• For one character: _</li> <li>• For any number of characters: %</li> </ul> <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"> <li>• Index name</li> <li>• Columns contained in index</li> <li>• Position of column in index</li> <li>• Specifies whether index is UNIQUE</li> <li>• Sort sequence</li> </ul>
\di [sconnect]	Logs the user off of the database
\dp [PATTERN]	<p>Lists all database procedures that correspond to the specified [PATTERN].</p> <p>[PATTERN] can be specified as follows: [SCHEMA.] [OBJECT_NAME]. The following placeholders are possible:</p> <ul style="list-style-type: none"> <li>• For one character: _</li> <li>• For any number of characters: %</li> </ul> <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"> <li>• Schema name</li> <li>• Name of the database procedure</li> <li>• Package to which database procedure is assigned</li> </ul>

Command	Description
\ds [NAME]	<p>Lists all schemas that correspond to the specified [NAME].</p> <p>[NAME] can be specified as follows: [SCHEMA.] [OBJECT_NAME]. The following placeholders are possible:</p> <ul style="list-style-type: none"> <li>• For one character: _</li> <li>• For any number of characters: %</li> </ul> <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"> <li>• Schema Name</li> <li>• Owner</li> </ul>
\dt [PATTERN]	<p>Lists all tables that correspond to the specified [PATTERN].</p> <p>[PATTERN] can be specified as follows: [SCHEMA.] [OBJECT_NAME]. The following placeholders are possible:</p> <ul style="list-style-type: none"> <li>• For one character: _</li> <li>• For any number of characters: %</li> </ul> <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"> <li>• Schema name</li> <li>• Table name</li> <li>• Table type</li> </ul>
\du [NAME]	<p>Lists all database users that correspond to the specified [NAME].</p> <p>[NAME] can be specified as follows: [SCHEMA.] [OBJECT_NAME]. The following placeholders are possible:</p> <ul style="list-style-type: none"> <li>• For one character: _</li> <li>• For any number of characters: %</li> </ul> <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"> <li>• Name of the database user</li> <li>• User properties</li> </ul>

Command	Description
<code>\dv [PATTERN]</code>	<p>Lists all views that correspond to the specified [PATTERN] .</p> <p>[PATTERN] can be specified as follows: [SCHEMA.] [OBJECT_NAME]. The following placeholders are possible:</p> <ul style="list-style-type: none"> <li>• For one character: _</li> <li>• For any number of characters: %</li> </ul> <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"> <li>• Schema name</li> <li>• View name</li> <li>• View types</li> </ul>
<code>\e[dit] [&lt;file&gt;]</code>	Writes the command buffer to the specified file where you can edit it with an editor
<code>\f[ieldsep] &lt;separator&gt;</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>\i[nput] &lt;file&gt;</code>	Imports commands from the specified batch file
<code>\m[ode] &lt;INTERNAL SAPR3&gt;</code>	Changes the SQL mode
<code>\mu[ltiline] ON   OFF</code>	Switches multiple line mode on/off
<code>\o[utput] &lt;file&gt;</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>\q[uit]</code>	Exits HDBSQL
<code>\r[eset]</code>	Deletes the current command buffer
<code>\ro[wsep] &lt;separator&gt;</code>	Uses the specified separator character to separate the individual rows of the result
<code>\s[tatus]</code>	Displays general information about the database

## 10.2 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 &lt;host&gt; -i &lt;instance&gt; -u &lt;database_user&gt; -p &lt;database_user_password&gt;</pre>
Log onto a database in a multitenant database container	Run the following command all on one line: <pre>hdbsql -n &lt;host&gt; -i &lt;instance&gt; -u &lt;database_user&gt; -p &lt;database_user_password&gt; -d &lt;database_name&gt;</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 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>\c -n &lt;host&gt; -i &lt;instance&gt; -u &lt;database_user&gt; -p &lt;database_user_password&gt;</pre>

Option	Action
Log onto a database in a multitenant database container	Run the following command all on one line: <pre>\c -n &lt;host&gt; -i &lt;instance&gt; -u &lt;database_user&gt; -p &lt;database_user_password&gt; -d &lt;database_name&gt;</pre>

### 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
```

## 10.3 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 (;).

### 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
```

## 10.4 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 .
  - 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`

---

## 10.5 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.

## 10.6 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

Export a list of all schemas and all entries in the table HOTEL.CUSTOMER to an external file.

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 the following command: \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|

```

---

# Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities

SAP HANA server software and tools can be used for several SAP HANA platform and options scenarios as well as the respective capabilities used in these scenarios. The availability of these is based on the available SAP HANA licenses and the SAP HANA landscape, including the type and version of the back-end systems the SAP HANA administration and development tools are connected to. There are several types of licenses available for SAP HANA. Depending on your SAP HANA installation license type, some of the features and tools described in the SAP HANA platform documentation may only be available in the SAP HANA options and capabilities, which may be released independently of an SAP HANA Platform Support Package Stack (SPS). Although various features included in SAP HANA options and capabilities are cited in the SAP HANA platform documentation, each SAP HANA edition governs the options and capabilities available. Based on this, customers do not necessarily have the right to use features included in SAP HANA options and capabilities. For customers to whom these license restrictions apply, the use of features included in SAP HANA options and capabilities in a production system requires purchasing the corresponding software license(s) from SAP. The documentation for the SAP HANA options is available in SAP Help Portal. If you have additional questions about what your particular license provides, or wish to discuss licensing features available in SAP HANA options, please contact your SAP account team representative.

---

# Important Disclaimers and Legal Information

## Coding Samples

Any software coding and/or code lines / strings ("Code") included in this documentation are only examples and are not intended to be used in a productive system environment. The Code is only intended to better explain and visualize the syntax and phrasing rules of certain coding. SAP does not warrant the correctness and completeness of the Code given herein, and SAP shall not be liable for errors or damages caused by the usage of the Code, unless damages were caused by SAP intentionally or by SAP's gross negligence.

## Gender-Neutral Language

As far as possible, SAP documentation is gender neutral. Depending on the context, the reader is addressed directly with "you", or a gender-neutral noun (such as "sales person" or "working days") is used. If when referring to members of both sexes, however, the third-person singular cannot be avoided or a gender-neutral noun does not exist, SAP reserves the right to use the masculine form of the noun and pronoun. This is to ensure that the documentation remains comprehensible.

## Internet Hyperlinks

The SAP documentation may contain hyperlinks to the Internet. These hyperlinks are intended to serve as a hint about where to find related information. SAP does not warrant the availability and correctness of this related information or the ability of this information to serve a particular purpose. SAP shall not be liable for any damages caused by the use of related information unless damages have been caused by SAP's gross negligence or willful misconduct. All links are categorized for transparency (see: <https://help.sap.com/viewer/disclaimer>).



[go.sap.com/registration/  
contact.html](https://go.sap.com/registration/contact.html)

© 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/corporate/en/legal/copyright.html> for additional trademark information and notices.

**SAP**