

PUBLIC
SAP HANA Cockpit 2.0 SP 12
Document Version: 1.0 – 2020 -06-08

SAP HANA Administration with SAP HANA Cockpit

© 2020 SAP SE or an SAP affiliate company. All rights reserved.

THE BEST RUN SAP

Content

1 SAP HANA Administration with SAP HANA Cockpit. 8

2 Getting Started with SAP HANA Cockpit. 10

2.1 Set Up SAP HANA Cockpit for the First Time. 10

Determine Ports for SAP HANA Cockpit and Cockpit Manager. 11

Open the SAP HANA Cockpit. 13

2.2 The SAP HANA Cockpit Home Page. 14

2.3 The Database Directory Page. 14

Open the Database Directory. 14

2.4 The Database Overview Page. 15

2.5 The SAP HANA Cockpit Shell Bar. 16

Buttons on the Shell Bar. 16

Refreshing Data in the SAP HANA Cockpit. 17

Search, Sort, and Filter Buttons. 18

2.6 Personalize the SAP HANA Cockpit. 19

Personalizing the *Home* Page. 19

Personalizing the *Database Overview* Page. 21

Adjusting SAP HANA Cockpit User Settings. 22

Saved Views. 22

Emailing a Link to a Page. 23

2.7 Opening the SQL Console. 24

2.8 Security Considerations for SAP HANA Cockpit. 24

Setting Up Single Sign-On. 26

Importing a Certificate for Encrypted Communication. 32

Ensuring a Secure Browser Connection. 33

Data Protection and Privacy in SAP HANA Cockpit. 34

Auditing in SAP HANA Cockpit. 36

3 Setup and Administration with the Cockpit Manager. 37

3.1 Open the Cockpit Manager. 38

3.2 Cockpit Access. 38

3.3 SAP HANA Cockpit User Management. 39

Create or Enable an SAP HANA Cockpit User. 40

Add or Remove Database Groups for Cockpit Users. 42

Editing Settings for a Cockpit User. 43

Deactivate or Activate a Cockpit User. 44

Delete a Cockpit User or Revoke Cockpit Access. 45

3.4 Registered Databases. 47
The Technical User. 47
Register a Database. 48
Edit Database Settings. 52
Unregister a Database. 56
Export Databases. 56
Import Databases. 58
3.5 Database Groups. 61
Creating a Database Group. 62
Deleting a Database Group. 63
Adding or Removing Databases in a Database Group. 64
Adding or Removing Cockpit Users in Database Groups. 65
3.6 Managing Databases, Users, and Groups with the Cockpit APIs. 66
Using Cockpit POST APIs. 68
Using Cockpit GET APIs. 71
3.7 Cockpit Settings. 80
Setting Data Collection 80
Set Up a Proxy Server. 82
Setting Connection Timeout. 82
Displaying Auto-Generated Database Groups. 83
Enabling SSO with Kerberos. 84
3.8 User Notifications and Session Monitoring. 85
Monitoring Active Browser Sessions. 86
Sending a Notification to Logged-In Users. 87
3.9 View Logs to Troubleshoot the Cockpit. 88
3.10 Using XS CLI Commands to Troubleshoot the Cockpit. 89

4 Landscape Monitoring and Administration. 92
4.1 Monitoring Your Databases in the Database Directory. 92
Connect to a Database With SSO or SAP HANA Credentials. 97
4.2 Configurations and Configuration Templates. 98
Taking a Snapshot of a Database's Configuration. 99
Comparing Database Configurations. 99
Creating a Configuration Template. 100
Applying a Configuration Template. 101
Modifying a Configuration Template. 103
Deleting a Configuration Template. 103
4.3 SAP EarlyWatch Alert Service. 104
Specifying Authorization for the Technical User. 104
Managing SAP EarlyWatch Alert Global and Database-Specific Settings. 105
Viewing SAP EarlyWatch Alerts. 107

5 Using the Database Overview Page to Manage a Database. 109
5.1 Authorizations Needed for Monitoring and Administration. 111
5.2 Cards Available on the Database Overview Page. 119

6 Database Default Views. 124
6.1 Getting to the Database Overview Page. 125

7 Monitoring View. 126
7.1 Alerts. 126
Alert Summary and Details. 127
View Past Alerts. 129
Alert Configuration. 130
7.2 Memory Usage. 139
Memory Analysis. 139
Managing Warm Data With the Native Storage Extension. 147
View and Configure the Persistent Memory Monitor. 153
7.3 Services. 159
Service Details. 160
Operations on Services. 163
Start a Database. 164
Stop a Database. 166
7.4 System Replication. 167
Monitoring SAP HANA System Replication. 167
General Prerequisites for Configuring SAP HANA System Replication. 173
Configure SAP HANA System Replication. 176
Perform a Takeover. 182
Perform a Failback. 183
Disable SAP HANA System Replication. 185
7.5 Admission Control. 186
Manage Admission Control. 186
Manage Admission Control Advanced Analytics 187
Admission Control Default Configuration Parameters. 188
7.6 Monitoring Multi-Host Systems. 191
Monitoring the Network Between Multiple Hosts 191
Monitoring Multi-Host Health. 191
7.7 Disk Usage: Monitor Disk Volume. 194
Reclaim Space. 196

8 Security and User Management View. 197
8.1 Security. 197
SAP HANA Security Checklists and Recommendations. 197
SAP HANA Database Checklists and Recommendations. 199

Monitoring Critical Security Settings . . . 221
Managing Server-Side Data Encryption. . . . 229
Auditing Database Activity. . . . 253
Configure the Database Password Policy and Password Blacklist. . . . 274
Configure a Password Policy for a User Group. . . . 283
Managing Client Certificates. . . . 284
Data Anonymization. . . . 307
Display Information about an "Insufficient Privilege" Error. . . . 309
8.2 User and Role Management. . . . 310
SAP HANA Database Users vs Cockpit Users. . . . 311
View a Database User. . . . 312
View a Database Role. . . . 318
Create a Role. . . . 323
Change a Role. . . . 338
Delete a Role. . . . 339
Create a User Group. . . . 340
Create a Database User. . . . 349
Create a Restricted Database User. . . . 354
Assign Roles to Database Users. . . . 358
Assign Privileges to a User. . . . 360
Change a Database User. . . . 361
Deactivate a Database User. . . . 362
Delete a Database User. . . . 363
Add a SAML Identity Provider in SAP HANA Cockpit. . . . 365
Add a JWT Identity Provider in SAP HANA Cockpit. . . . 367
Resolve Object Authorization Errors. . . . 368

9 Administration View. . . . 371
9.1 Managing Tenants from an SAP HANA System Database. . . . 371
Assign the OS User and Group for High Isolation. . . . 372
Clear the OS User and Group When Decreasing Isolation. . . . 378
Create a Tenant Database. . . . 381
Start a Tenant Database. . . . 383
Stop a Tenant Database. . . . 384
Prevent the Start of a Tenant Database at System Startup. . . . 385
Rename a Tenant Database. . . . 386
Delete a Tenant Database. . . . 387
Restrict Features Available to a Tenant Database. . . . 389
Lock Parameters Against Editing for a Tenant Database. . . . 391
Create a Fallback Snapshot. . . . 397
Reset to a Fallback Snapshot. . . . 398
Delete a Fallback Snapshot. . . . 399

Copy or Move a Tenant Database Using Replication. . . . 400
Reset the SYSTEM Password of a Tenant Using the Cockpit. . . . 402
Memory and CPU Usage for Tenant Databases. . . . 403
Monitoring Tenant Databases in SAP HANA Cockpit. . . . 418
Add or Remove Services in a Tenant Database. . . . 422
Start and Stop Services in a Tenant Database. . . . 425
Change the Port of a Service in a Tenant Database. . . . 426
9.2 Database Administration. . . . 427
Configure Host Failover. . . . 427
License Management. . . . 433
Database Configuration in SAP HANA Cockpit. . . . 443
Managing Workload Classes. . . . 457
9.3 Recommendations. . . . 470
View and Follow Recommendations. . . . 471
Configure Recommendations. . . . 472
Use the NSE Advisor for Warm Data Recommendations. . . . 473
9.4 Alerting and Diagnostics. . . . 475
Collect and Download Diagnosis Information. . . . 475
9.5 Database Information. . . . 477
9.6 Manage Hadoop Clusters. . . . 478
9.7 SAP HANA Smart Data Access. . . . 478
Monitor Remote Statements and Connections. . . . 479
9.8 Table Distribution. . . . 481
Table Placement Rules. . . . 481
Designing Partitions. . . . 486
View or Modify a Table Distribution. . . . 489
View Redistribution Execution History. . . . 502
Generate and Execute a Table Redistribution Plan 506
Table Group Advisor. . . . 509
Table Replication. . . . 510

10 Monitoring, Analyzing, and Improving Performance. . . . 513
10.1 Monitoring Performance in SAP HANA Cockpit. . . . 514
The Performance Monitor. . . . 515
Threads. . . . 521
Session Monitoring. . . . 525
Statements. . . . 525
Blocked Transactions. . . . 530
Monitor Table Usage. . . . 530
10.2 Analyzing Performance in SAP HANA Cockpit. . . . 531
Capturing and Replaying Workloads. . . . 532
Analyzing Workloads. . . . 561

Analyzing Statement Performance. . . . 568
Analyzing Memory with the Memory Profiler. . . . 579
10.3 Improving Performance in SAP HANA Cockpit. . . . 584
Data Cache. . . . 585
SQL Plan Stability. . . . 588
Statement Hints. . . . 590

11 Backup and Recovery. . . . 591
11.1 Display the Backup Configuration Settings. . . . 592
11.2 Change the Backup Configuration Settings. . . . 592
Backup Configuration Settings. . . . 594
11.3 Create Data Backups. . . . 603
11.4 Create a Data Snapshot (Native SQL). . . . 606
11.5 Schedule Backups. . . . 611
Manually Enable the Data Backup Scheduler. . . . 616
Manage Backup Schedules. . . . 619
11.6 Cancel a Backup. . . . 620
11.7 Housekeeping: Deleting and Archiving Backups. . . . 621
Prerequisites for Deleting Old Backups. . . . 622
Delete Backup Generations. . . . 623
Delete Backups. . . . 625
11.8 Recover a Database. . . . 628
Recover SAP HANA From a Data Snapshot. . . . 632
11.9 Cancel a Recovery. . . . 635
11.10 Resume a Canceled Recovery. . . . 637
Resuming a Canceled Recovery Using Native SQL. . . . 639
11.11 Copy a Database. . . . 640

1 SAP HANA Administration with SAP HANA Cockpit

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

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

What can I do with the cockpit?

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

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

How can I keep an eye on the big picture?

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

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

How do I get access to groups of databases?

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

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

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

About SAP HANA Administration with SAP HANA Cockpit

This documentation provides information on the Cockpit Manager configuration tool, and on all the features available through the aggregate and system overviews within the SAP HANA cockpit. The SAP HANA database explorer and other linked applications such as the SAP HANA database lifecycle manager and the SAP HANA XS advanced cockpit are documented separately. Also refer to the *SAP HANA Administration Guide* as the comprehensive source of information for administering SAP HANA using all SAP HANA administration tools.

Related Information

SAP Note 2380291

2 Getting Started with SAP HANA Cockpit

Use the SAP HANA cockpit to administer and monitor your HANA database.

2.1 Set Up SAP HANA Cockpit for the First Time

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

Prerequisites

- You have access to the cockpit administrator (COCKPIT_ADMIN) user, created during the installation process
- You have access to the master password that you were prompted to enter during the installation process.
- You know the URL for the cockpit created during the installation process.

Procedure

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

   Access the Cockpit Manager by clicking *Manage Cockpit* on the cockpit home page.
2. Create other cockpit users, keeping in mind what cockpit actions you want them to be able to perform.
3. Register databases.
4. Create database groups and add registered databases to each group.
5. Assign cockpit users access to one or more groups of databases.
6. Determine whether you would like to configure single sign-on authentication to the cockpit and/or the registered databases.
7. Share the credentials of the newly created cockpit users with the appropriate people, and instruct them to sign in to the SAP HANA cockpit.

   In the cockpit, the *Database Directory* displays all the registered databases to which this cockpit user has access.

Related Information

Open the SAP HANA Cockpit [page 13]
Registered Databases [page 47]
Setting Up Single Sign-On [page 26]
SAP HANA Cockpit User Management [page 39]
SAP Note 2380291
SAP Note 2535229
SAP Note 2594513

2.1.1 Determine Ports for SAP HANA Cockpit and Cockpit Manager

The ports for SAP HANA cockpit and the Cockpit Manager can be determined in the XS console after the cockpit installation.

Prerequisites

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

Context

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

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

Procedure

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

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

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

2. Log on to the SAP HANA XS advanced runtime by using the following command:

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

3. Enter the password for the XSA admin user.
4. Display a list of the applications running in the current space. In the command shell, run the following command:

```
xs apps
```

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

Output Code

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

```
cockpit-xsa-svc            STARTED          1/1          768 MB
<unlimited>   https://<hostname>:51024
cockpit-admin-web-app      STARTED          1/1          128 MB
<unlimited>   https://<hostname>:51023
cockpit-web-app            STARTED          1/1          512 MB
<unlimited>   https://<hostname>:51021
```

Related Information

SAP Note 2389709 - Specifying the port for SAP HANA Cockpit before installation

2.1.2 Open the SAP HANA Cockpit

Access the SAP HANA cockpit from a web browser.

Prerequisites

- You know the URL for the cockpit that was created during the cockpit installation process.
- You have a cockpit user name and password.
- Your Web browser supports the SAPUI5 library `sap.m`.For more information about SAPUI5 browser support, see SAP Note 1716423. See also the *Product Availability Matrix (PAM) for SAPUI5*.

Procedure

1. Enter the SAP HANA cockpit URL in your browser.
2. Enter your cockpit user name and password.
   SAP HANA cockpit opens.

Related Information

Security Considerations for SAP HANA Cockpit [page 24]

SAP Note 1716423

Product Availability Matrix (PAM) for SAPUI5

2.2 The SAP HANA Cockpit Home Page

The SAP HANA cockpit home page appears once your logon credentials are validated.

The SAP HANA cockpit home page contains multiple tiles, organized into sections. The *Monitor Landscape* section contains tiles that let you access registered databases. Only those groups and databases to which the logged on user has been granted permissions (using SAP HANA Cockpit Manager) appear. The same database can appear in multiple groups, but it doesn't matter which group you use to access the database.

The *Manage Landscape* section contains tiles that let you manage the SAP HANA cockpit, browse database objects, or execute SQL statements using SAP HANA Database Explorer.

The *My Home* section, if available, contains short cut tiles to applications within a registered database.

2.3 The Database Directory Page

The *Database Directory* page contains information about all the registered databases to which you have been granted access.

For each tenant database listed, you can see information on the status of the database and alerts. For system databases, you can see information on memory, CPU, and disk use.

Access to the information for both tenant and system databases is based on the currently cached credentials, which can also be changed form the *Database Directory* page.

For more information on using the *Database Directory* page, see Landscape Monitoring and Administration [page 92].

2.3.1 Open the Database Directory

Display a list of registered databases to which you have been granted access.

Context

Depending on how your SAP HANA cockpit *Home* page is configured, there are several ways to access the *Database Directory*.

Procedure

1. On the SAP HANA cockpit *Home* page, select one of the following:

- The *Database Directory* tile
- The auto-generated tile for the system type (Production, Development, Test) of the registered database
- The *Group* tile to which the registered database is assigned

2. (Optional) Select a database from the *Database Directory* list to open the *Database Overview* for that database, or right-click a database to open it in a new tab or window.

Related Information

Search, Sort, and Filter Buttons [page 18]

2.4 The Database Overview Page

Learn how to navigate and use the *Database Overview* page.

The *Database Overview* page contains a series of cards. Each card provides a starting point to a group of related applications that allow you to monitor and manage the database.

The cards are grouped into four views:

- Monitoring
- Security and User Management
- Administration
- All

To change the view, at the top left of the page, click the ⌄ icon beside the current view name. Cockpit remembers the last view selected and automatically loads it the next time you open SAP HANA Cockpit. You can rearrange cards in a view by dragging a card to a new location, but you can't add or remove cards to or from the view.

When moving your mouse over a card, areas that change color and display a hand ( ) indicate an available link. For example, on the *Memory Usage* card, a hand and a color change appear when you move over the card title, *Used Memory*, *Resident Memory*, and *Monitor Performance* or within the card. On the *Monitoring* card, nothing happens when you hover over the title, but the hand and color change happens when you hoover over each item within the card. Some cards provide summary information about the database. This information is collected from the card's underlying applications. Depending where you click on these cards, you can change the summary information displayed, open the card's application to display filtered details, or open a related application from another card.

To find the links on a card, move your mouse over the card. Areas that change color indicate an available link. For example, on the *Memory Usage* card there are four areas that contain links.

- 1 and 4 - Opens the Performance Monitor application.
- 2 and 3 - Lets you change the summary memory information displayed on the card.
- 5 - Displays links to other memory related-applications.

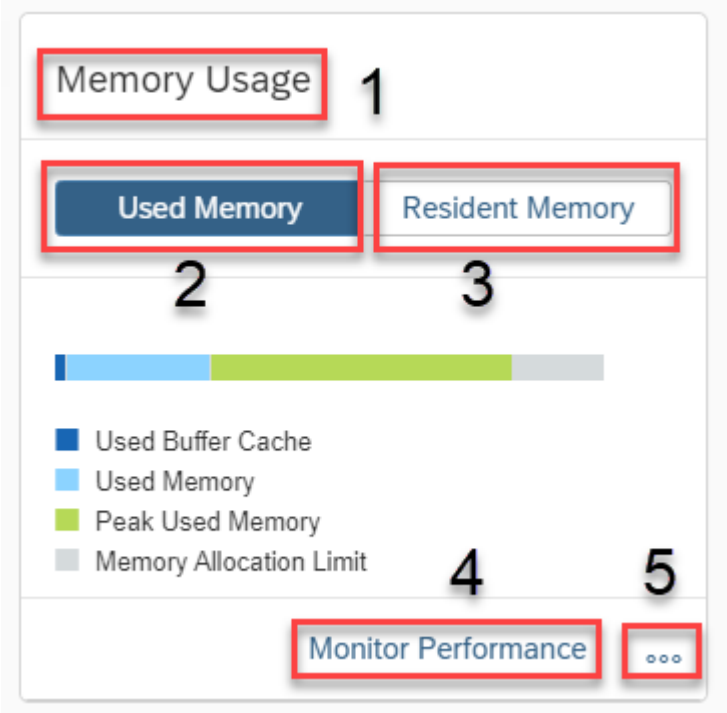

When connected to the system database (SYSTEMDB), buttons to stop or start the system appear at the top of the page. When the system is running, a link to the *Database Management* application also appears on the page.

2.5 The SAP HANA Cockpit Shell Bar

There are several buttons on the SAP HANA Cockpit shell bar. Some are appear on the top of every page while others appear only on relevant pages.

2.5.1 Buttons on the Shell Bar

A shell bar appears at the top of every page that contains a common set of buttons.

There are several buttons on the left side of the header.



The *<cockpit_user>* button lets you customize some cockpit settings, manage the cards on the *Database Overview* page, or sign out of SAP HANA Cockpit.

The *(Back)* button takes you back one screen with each click.

> ⚠ Caution
>
> Don't use the browser *Back* button as it can result in unexpected behavior.

The *(Home)* button takes you directly to the cockpit *Home* page from any application page in cockpit.

The tenant database name and current database user name (for example HA2@HA2 (SYSTEM)) displays status and general information about the current database when you click it. If the status indicates issues, you can't access details on the issues from there. You need to open the *Manage Services* application either by using the *Services* card on the *Database Overview* page or by using *Search for Applications* on the *Navigation menu*, which is in the top, middle of the shell bar.

In the middle of the shell bar, beside the page name, is the *(Navigation menu)* button. Every *Navigation menu* contains links to the *Database Overview*, *Database Directory*, and *Home* pages. Where applicable, it suggests links to applications related to the current application page. If the application link you want isn't displayed, use *Search for Applications* to find the link. You can filter the application list or click the button to display the full list.

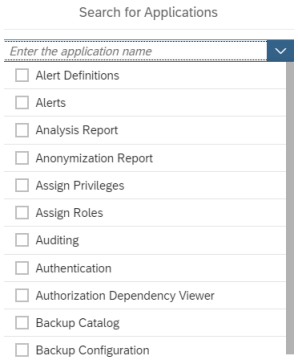

To select a single application, click the application name. To select multiple applications, click the checkbox. When you have finished selecting applications, click the button to close the list. Click *Open*. Each selected application opens in a separate tab in the browser.

There are four buttons on the right side of the header.



You can change the database by clicking the *(Switch Database)* button. Only those databases you have permissions to are listed.

You can refresh the page by clicking the *(Refresh)* button.

You can schedule the frequency with which the display is refreshed by clicking the *(Schedule Refresh)* button.

Finally, you can access the online help by clicking the button.

2.5.2 Refreshing Data in the SAP HANA Cockpit

Schedule how frequently data on a page is updated.

By default, the data displayed on a page within the SAP HANA cockpit doesn't refresh automatically, but you can change this behavior by setting an interval to automatically refresh the screen. Even when a refresh interval is active, you can still refresh the screen on demand. Once the refresh interval is active, it applies to all pages

within the SAP HANA cockpit, not just the page you set it on. The refresh interval persists between pages until the active SAP HANA cockpit user logs out or closes the browser.

To refresh the display on demand, click the *(Refresh)* button.

To activate a refresh interval, click the *(Schedule Refresh)* button and then select an interval. You can't add intervals to the list when the refresh interval is active. When active, the refresh button changes to ⊙. You can't check the frequency when refresh is active. You must stop the refresh and then reactivate it.

To stop the refresh interval, click ⊙. The button changes to *(Schedule Refresh)*.

2.5.3 Search, Sort, and Filter Buttons

Use these buttons to work with tables of data.

> **i Note**
>
> Not all pages support all buttons. When supported, the indicated button appears.

Filter	Filtering reduces the displayed rows to contain the specified values in the specified columns. If you select multiple values on a single list, then the values are treated as OR operations. But if you define filters on multiple columns, then the columns are treated as AND operations. To filter data by column, select one or more values from one or more predefined dropdown lists. Use the *(Pin header on press)* button to keep the filter fields on the page when scrolling page through multiple screens of data. Not all columns in the data table can be filtered and you can't add new column filters. You can save filters as views. See *Using My Views*.
Search	Searching reduces the displayed rows to contain the specified text string in any column. Searching uses exact match only. To search, begin typing a text string in the *(Search)* field. The displayed rows are instantly reduced by the string. You don't need to press *Enter* to apply the search. Searching can be used along with filtering. Search criteria isn't saved as part of a view. See *Using My Views*.
Hide/Show Filter Columns	You can choose which filter columns to display. Hidden columns remain hidden until you leave the current page. Hiding a filter column doesn't hide the corresponding column from the data table and any filter values defined on hidden columns are still applied to the data table. When you save filters as a view, hidden filter columns are automatically redisplayed in the saved view. To hide filter columns, click *Adapt Filters* and deselect a column to hide it.

Change sorting rules	You can modify the sorting rules to: • List data in descending or ascending order • Sort by or Group by object. Click the *(Group)* button. Choose the rules and click *OK*.
Table Personalization	You can choose which data column to display in the data table. You can hide/display columns in the data table using the *(Table Personalization)* button. table personalizations are visible to all database users for the current database, for the logged on cockpit user. Customizations persist until the cockpit user logs out. Upon the next logon, all customizations to the data table revert to the default values.

Related Information

Saved Views [page 22]

2.6 Personalize the SAP HANA Cockpit

Customize the appearance of the SAP HANA cockpit.

Since the cockpit is role-based, only apps and tiles relating to your role are displayed. Some customizations are tied to the active SAP HANA cockpit user, while others are tied to both the SAP HANA cockpit user and registered database.

2.6.1 Personalizing the *Home* Page

You can customize the SAP HANA cockpitdefault *Home* page.

Changes to the default *Home* page are specific to the active cockpit user. Changes are instantly applied; no save required. Be careful what you change, since there's also no forget-my-changes option when leaving the page.

Existing tiles can be rearranged directly on the *Home* page by dragging a tile within a group or between groups. All other customization actions require you to click the *<user_name>* icon and then click *Edit Home Page*.

The *My Home* group only appears as long as at least one tile has been added to the group. It continues to appear until you remove the last tile from the group. To add the first tile to *My Home*, see *Editing the Home Page*.

The *Monitor Landscape* and *Manage Landscape* groups are auto-generated and can be hidden. See *Displaying Auto-Generated Database Groups* in Related Information.

Related Information

Displaying Auto-Generated Database Groups [page 83]
Editing the Home Page [page 20]

2.6.1.1 Creating a Shortcut Tile

Add a shortcut tile to a page on the *Home* page.

Context

Create a shortcut tile on the *Home* page to take you directly to an application page.

> **i** Note
>
> Not all pages support shortcut tiles, but when it does, then the *(Save a Tile/Send Email)* icon appears in the top-right corner of the page.

The shortcut tile is specific to the logged on cockpit user and the registered database. If you click a tile and the current database credentials are insufficient, no data is loaded.

Shortcut tiles only appear on the *Home*.

Procedure

1. Go to the page you want to create a shortcut to.
2. Click the *(Save a Tile/Send Email)* icon.
3. Select *Save as Tile*.
4. (Optional) If prompted, then give the tile a name and specify the group for the tile. Otherwise, the tile is added to the *My Home* group with the name of the application.

2.6.1.2 Editing the Home Page

Customize the appearance of the *Home* page.

Changes to the *Home* page are tied to the active SAP HANA cockpit user. You can rearrange existing tiles directly on the *Home* page by dragging a tile within a group or between groups.

For all other customizations, click the *(<cockpit_user>)* icon and then select *Edit Home Page*.

To add a tile to the *My Home* section, drag a tile to the section.

To rename a section (other than *My Home*), click the title, and type the new name. To move a section, left-click and drag the section title to the new location.

You can only add shortcut tiles to the *Home* page. See *Create a Shortcut Tile* in Related Information.

Only tiles with the ⊗ *(Delete)* icon in its top-right corner can be deleted.

You can't:

- Rename or move the *My Home* section.
- Create new sections.
- Hide or rename tiles.
- Create new tiles (not shortcuts)

For sections with the *Reset* button, *Reset* reverts the section to its default name and tiles. Tiles moved to the section from another section disappear. They don't reappear in the original section. To preserve these moved tiles, before resetting the section, move them temporarily to another section and then move them back after the reset. If you forgot to move them before resetting, to restore the tiles, reset its original section.

If you reset a section in which you moved tiles to another section, the moved tiles appear in both sections. To remove duplicate tiles, drag the duplicate tile back to its original section, and reset the original section. *Reset* doesn't restore a moved section to its original location.

For the *Monitor Landscape* section, you can choose which auto-generated database type tiles appear. See *Displaying Auto-generated Database Groups* in Related Information.

Related Information

Displaying Auto-Generated Database Groups [page 83]
Creating a Shortcut Tile [page 20]

2.6.2 Personalizing the *Database Overview* Page

You can customize the *Database Overview* page of each system or tenant database.

When you access the *Database Overview* page of a tenant or system database, you have access to many cards, each representing a specific SAP HANA cockpit application. (If you don't have the required privileges, specific cards, features, or actions aren't available to you.) You can drag and drop to arrange these cards in your preferred order. You can filter the cards using one of four default views or hide cards that aren't relevant to you.

To hide cards, on the *Database Overview* page, select the view that contains the cards to be hidden. Click the 👤 *(<cockpit_user>)* icon, then *Manage Cards*. Only the cards for the select view are listed. Hide or display cards as desired. Click *OK* to save the settings.

While most actions in the 👤 *(<cockpit_user>)* area are available independently of the current context, some of the actions are directly tied to the content shown in the main content area, including:

- Settings
- Tools to personalize the current content

- Recently used and frequently visited apps

In the *(<cockpit_user>)* area, if Single Sign-On authentication is not enabled for your SAP HANA cockpit user, you can also change your SAP HANA cockpit password by selecting *User Account* and *Change Password* within *Settings*.

Related Information

Using the Database Overview Page to Manage a Database [page 109]
Monitoring Your Databases in the Database Directory [page 92]

2.6.3 Adjusting SAP HANA Cockpit User Settings

You can set preferences for the current SAP HANA cockpit user.

Changes to user settings are tied to the active SAP HANA cockpit user. To begin customizations, click the *(<cockpit_user>)* icon on the *Home* page, and select *(Settings)*.

If Single Sign-On authentication is not enabled for your user, then you can change the current SAP HANA cockpit password by selecting *User Account* and *Change Password*.

User Activities lets you display lists of recently and frequently user applications.

2.6.4 Saved Views

Save a set of filters in a view for reuse later.

Context

Views are a quick way to reuse a saved set of filters on a page. Views are cockpit user and database specific. They're also specific to the page they're created on. You can't copy views between pages, databases, or cockpit users.

Views are available on any page that allows filtering except the *Database Overview* page. The name of the current view appears in the top left of the page. The system view, *Default*, can't be modified, but changes to the default view can be saved as a new view.

Cockpit automatically loads the last view used for the page, for the logged on cockpit user. Always check the name of the loaded view and the defined filters to ensure that you understand the displayed data.

Procedure

1. To display the saved views for the page, click the ⌄ *(Select Views)* icon beside the 🔍 *(Search)* field.
2. To use a view, click the name.

   The filters are applied and the page refreshes. You can modify the filters of a loaded view. An * appears beside the view name to indicate unsaved changes.

3. To save filter changes, click the ⌄ *(Select Views)* icon.
   a. To update the existing view, select *Save*.
   b. To create a new view, select *SaveAs* and give the view a name.

   > ⚠ Caution
   >
   > You can only save the changes to the current view or to a new view. If you try to apply the changes to another existing view by selecting it, then the existing view is loaded, and the filter changes are lost.

4. To delete a view, in *My Views*, select *Manage*. Click the ⊗ *(Delete View)* icon beside the view.

2.6.5 Emailing a Link to a Page

Generate an email link to a page.

Context

The link generated is to the live page, not a snapshot of the page. The URL uses the current database credentials currently registered for the database.

If the logged on cockpit user doesn't have access to the database, then an error appears. If the current database credentials don't have the required privilege to view the data, then no data appears.

Not all pages allow you to email a link, but if it does, then the ⎋ *(Save a Tile/Send Email)* icon appears in the top-right corner of the page. Shortcut tiles only appear on the *Home*.

Procedure

1. Go to the page you want to email the link for.
2. Click the ⎋ *(Save a Tile/Send Email)* icon.
3. Select *Send Email*.

2.7 Opening the SQL Console

Access the SQL console in SAP HANA database explorer.

Procedure

1. From the SAP HANA cockpit *Home* page, click *Execute SQL*.
2. From the *Database Overview*, click *Open SQL Console* .

Results

The SQL console tab opens in SAP HANA Database Explorer. You're connected to the SAP HANA instance with the same database credentials from SAP HANA Cockpit.

2.8 Security Considerations for SAP HANA Cockpit

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

User Authentication

Several types of credentials are used within SAP HANA cockpit:

Credential	Details
COCKPIT_ADMIN	The master user for the cockpit created during the installation process. The password for the cockpit administrator user is the master password established during the installation process. This master user is assigned all three administrator roles, and can therefore access all aspects of the Cockpit Manager, and can create users, register databases, and assign users and databases to database groups.
Cockpit Users	The business users with access to the cockpit. Each is assigned one or more roles. For details, see *Managing Cockpit Users* in Related Information.

Credential	Details
Technical User	An application global per-database set of database credentials for access to remote databases and necessary to regularly gather information such as state, status, and other generalized KPIs. The technical user is a dedicated database user. It's required to register a database and can be created during database registration if it doesn't already exist. At a minimum, the technical user requires `CATALOG READ` system privilege and `SELECT` permission on `_SYS_STATISTICS` catalog. Human users should never log in using the technical user account.
Database User (User Remote Logon)	The per database/per user set of database credentials for a remote database. This user is used by the cockpit user to view more sensitive information, and to make changes within their roles as defined on that database. Each cockpit user needs to provide database user credentials with the system privilege CATALOG READ and SELECT on _SYS_STATISTICS in order to drill down in the cockpit to the overview information for a specific database. The cockpit securely encrypts and stores separate database credentials for each cockpit user, but you can clear and re-enter the credentials through the *Database Directory* when you desire to do so. We recommend that each cockpit user connects with different database user credentials.
Single Sign-On (SSO)	For any systems running version SAP HANA 2.0 SPS 01, or later, the cockpit database administrator has the option to enable or enforce SSO. See *Setting Up Single Sign-On* in Related Information.
Operating System User	A per database set of credentials for accessing the SAP Control process (starting and stopping the database, and restoring features). This user is usually the `<sid>adm` account. The cockpit securely encrypts and stores these credentials, but you can clear and re-enter the credentials through the *Database Directory* when you desire to do so.
Internal Communication	Service-to-service authentication
SAP HANA Service Broker User	For application persistence using the application's SAP HANA express database

Network and Communication Security

The cockpit uses secure protocols on all client browser connections to HTTPS ports. Communication to SAP HANA databases uses JDBC, and can be secured by importing certificates into the cockpit. Additional communication is made to the remote hosts using a restful interface, which also can be secured. You can use

properly signed certificates for the cockpit's external ports as well. For more information about obtaining certificates, see *Certificate Management in SAP HANA* in the *SAP HANA Security Guide.*

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

Related Information

SAP HANA Cockpit User Management [page 39]
Editing Settings for a Cockpit User [page 43]
Monitoring Your Databases in the Database Directory [page 92]
User and Role Management [page 310]
Importing a Certificate for Encrypted Communication [page 32]
Setting Up Single Sign-On [page 26]

2.8.1 Setting Up Single Sign-On

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

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

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

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

> **i Note**
>
> Before enabling SSO, consider migrating the Personal Security Environment (PSE) file to an in-database store. When SSO is enabled, a new PSE file may be created, which may prevent cockpit access to stored certificates. See *SAP Note 2656666*.

Related Information

Edit Database Settings [page 52]

Enforce Single Sign-On on the Database [page 54]
SAP Note 2656666

2.8.1.1 Configuring SSO Access to the SAP HANA Cockpit

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

Context

Authentication standard supported for SSO to SAP HANA cockpit: SAML

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

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

Procedure

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

      In the example above, we use cockpit roles for group names; you can choose your own group names.
   d. Select *Save*.
   e. Exit the SAP Cloud Platform Identity Authentication Administration Console.

2. Use the XS advanced cockpit to add a new SAML identity provider.
   a. Retrieve SAML2 metadata from your IDP.

      For example, using the SAP IDP, navigate to

      `https://<IDP_URL>/saml2/metadata`
   b. Use the `xs a` command to access the xsa-cockpit URL.
   c. In the left pane, select *Security* > *Trust Configuration*.
   d. Select *New Trust Configuration* to add a new provider.
   e. In the metadata text box, enter XML based on the SAML metadata you retrieved from the IDP in step 2.a [page 28].
   f. Select *Parse* to fill in the remaining fields.
   g. Select *Save*.
3. In the XS advanced cockpit, map the role collections. (For more information, see SAP Note 2569903).
   a. Select *Security* > *Trust Configuration.*
   b. Select the trust configuration you created in step 2 [page 28].
   c. Select *Role Collection Mappings*.
   d. Select *New Role Collection Mapping* to map each of the role collections to a group within the IDP. If you use the sample group names in step 1.c [page 27], the mappings are:

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

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

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

   `https://<cockpit_FQDN>3<instance#>32/uaa-security/saml/metadata`

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

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

d. Select *Save*.

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

> **i Note**
>
> Only one application can be configured with the SAML provider. An error occurs if you introduce a duplicate, even with a different name.

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

Related Information

Setting Up Single Sign-On [page 26]
SAP Note 2569903

2.8.1.2 Configuring SSO Access to a Database

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

Prerequisites

- The database has been registered in the cockpit. It meets these version restrictions:
  - SAP HANA 1.0 SPS 12 revision 14 or later, or
  - SAP HANA 2.0 SPS 01 or later
- You have a database user with the CATALOG READ, TRUST ADMIN, CERTIFICATE ADMIN, and USER ADMIN privileges.
  The database user we use in the following steps is SSO_USER.
  To assign the necessary system privileges to an existing user called SSO_USER, execute these SQL statements:

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

- You have a second database user with the USER ADMIN system privilege.
  The second database user we use in the following steps is USER_ADMIN.
- A user with the Cockpit Administrator or Cockpit User Administrator role has created one or more cockpit users who are accessing this database and assigned them the role `Cockpit User`.
  The cockpit user we use in the following steps is COCKPIT_USER.
- For step 1 [page 31], you have a cockpit user with the Cockpit Administrator or Cockpit Database Administrator role.
  The cockpit administrator we use in the following steps is COCKPIT_ADMIN.

Context

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

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

The following steps configure SSO for a single managed SAP HANA database. Perform them for each database where SSO is needed. There's a separate set of steps for configuring SSO for the cockpit itself—see Related Information.

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

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

Procedure

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

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

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

Next Steps

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

Related Information

Create or Enable an SAP HANA Cockpit User [page 40]
Configuring SSO Access to the SAP HANA Cockpit [page 27]
SAP Note 2656666

2.8.2 Importing a Certificate for Encrypted Communication

You can import a certificate to enable an encrypted HTTP connection from the cockpit to SAP Control, or an encrypted JDBC connection from the cockpit to an SAP HANA database.

Procedure

1. Sign on as `<sid>adm` to the remote system with which you want the cockpit to establish a connection.
2. On the remote system, run the `sapgenpse` tool to export certificates from the in-memory certificate collection, or from the file-system PSE (`$SECUDIR/sapgenpse export_own_cert`).

   The location of required certificates depends on how you manage certificates in your system. For instance, point to SAP Control from a browser, and obtain the certificate.
3. Sign on as `<sid>adm` to the system hosting the cockpit.
4. Run the command `XS login`, and provide the name and password of the `COCKPIT_ADMIN` so that the `<sid>adm` can execute the xs command-line tool in the user context of `COCKPIT_ADMIN`.
5. In SAP HANA XS advanced model (XSA), trust the certificates using the command syntax `xs trust-certificate <ALIAS> -c <CERT_FILE>`, where `<ALIAS>` is a unique alias for the certificate within XSA, for example. `BZ1_SAPCONTROL` and `<CERT_FILE>` is the certificate you exported from the remote system.
6. Run the commands `xs restage cockpit-hdb-service cockpit-hdb-svc` followed by `xs restart service cockpit-hdb-svc` to refresh the certificates for the cockpit.

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

Related Information

SAP Note 2300943
Register a Database [page 48]

2.8.3 Ensuring a Secure Browser Connection

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

Context

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

> **i Note**
>
> The steps outlined here are deliberately generic so as to apply to all configuration scenarios. For a detailed, specific example, see SAP Note 2631903.

Procedure

1. To use a certificate from a certificate authority:
   a. Obtain a signed certificate from the certificate authority.
   b. In SAP HANA XS advanced model (XSA), verify the domain using the command `xs domains`.
   c. Set the XS domain certificate using the command `xs set-certificate <domain> -c <CERT_FILE> -k <CERT_KEY>`.
   d. Restart XSA using the command `XSA restart`.
2. To use your own certificate:
   a. Create a self-signed root certificate.
   b. Create a second certificate that is signed by the root certificate
   c. In SAP HANA XS advanced model (XSA), verify the domain using the command `xs domains`.
   d. Set the XS domain certificate using the command `xs set-certificate <domain> -c <CERT_FILE> -k <CERT_KEY>`.
   e. Restart XSA using the command `XSA restart`.
   f. Update your browser's list of trusted root certificates to include the newly created root certificate file.

   > **i Note**
   >
   > The XS domain name is tied to the certificate request. If the domain name starts out basically as the host name, you can change the domain name to a virtual host name so it's easy to remember and easy to redirect. The certificate for the domain needs to contain both the actual and virtual host names.

Related Information

SAP Note 2631903
SAP Note 2243019
SAP Note 2300936

2.8.4 Data Protection and Privacy in SAP HANA Cockpit

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

Introduction

Data protection is associated with numerous legal requirements and privacy concerns. In addition to compliance with general data privacy regulation, it's necessary to consider compliance with industry-specific legislation in different countries/regions. SAP provides specific features and functions to support compliance with regard to relevant legal requirements, including data protection. SAP doesn't give any advice on whether these features and functions are the best method to support company, industry, regional, or country-specific requirements. Furthermore, don't take this information as advice or a recommendation regarding additional features that would be required in specific IT environments. Decisions related to data protection must be made on a case-by-case basis, considering the given system landscape and the applicable legal requirements.

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

Glossary

Term	Definition
Consent	The action of the data subject confirming that the usage of personal data can be allowed for a given purpose. A consent functionality allows the storage of a consent record in relation to a specific purpose and shows if a data subject has granted, withdrawn, or denied consent.
Deletion	The irreversible destruction of personal data.

Term	Definition
Personal data	Any information relating to an identified or identifiable natural person ("data subject"). An identifiable natural person is one who can be directly or indirectly identified. This identification can be by reference to an identifier such as a name, an identification number, location data, or online identifier. Personal data can also refer to physical, physiological, genetic, mental, economic, cultural, or social factors of that natural person.
Purpose	A legal, contractual, or in other form justified reason for the processing of personal data. The assumption is that any purpose has an end that is already defined when the purpose starts.

User Consent

SAP HANA cockpit stores only personal data entered by users; it only collects personal data with a user's knowledge.

Logging Read Access and Changes

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

Information Report

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

Deletion of Personal Data

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

- *Edit Database Settings*
- *Delete a Cockpit User or Revoke Cockpit Access*

Related Information

Auditing Database Activity [page 253]

Edit Database Settings [page 52]
Delete a Cockpit User or Revoke Cockpit Access [page 45]

2.8.5 Auditing in SAP HANA Cockpit

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

Prerequisites

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

Context

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

Procedure

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

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

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

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

Related Information

SAP HANA Administration Guide for SAP HANA Platform

3 Setup and Administration with the Cockpit Manager

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

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

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

- Create and manage cockpit users

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

- Register databases
- Create database groups
- Add databases to database groups
- Grant cockpit users access to database groups

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

- Modify settings related to the configuration of the cockpit.

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

Related Information

Set Up SAP HANA Cockpit for the First Time [page 10]
Determine Ports for SAP HANA Cockpit and Cockpit Manager [page 11]
Open the SAP HANA Cockpit [page 13]
SAP HANA Cockpit User Management [page 39]
Registered Databases [page 47]
Cockpit Settings [page 80]
Security Considerations for SAP HANA Cockpit [page 24]
Using XS CLI Commands to Troubleshoot the Cockpit [page 89]

3.1 Open the Cockpit Manager

You access the SAP HANA Cockpit Manager from a web browser.

Prerequisites

- You know the URL for the cockpit manager, created during the cockpit installation process.
- You have a cockpit user name and password with sufficient roles to use SAP HANA Cockpit manager.
- Your Web browser supports the SAPUI5 library `sap.m`.For more information about SAPUI5 browser support, see SAP Note 1716423. See also the *Product Availability Matrix (PAM) for SAPUI5*.

Procedure

1. Enter the URL for SAP HANA Cockpit Manager or SAP HANA cockpit in your browser.
2. Enter your cockpit user name and password.
3. For SAP HANA cockpit only, click the *Manage Cockpit* tile under *Manage Landscape*.

Related Information

Security Considerations for SAP HANA Cockpit [page 24]
SAP Note 1716423
Product Availability Matrix (PAM) for SAPUI5

3.2 Cockpit Access

Access to SAP HANA cockpit is controlled using a combination of users, groups, and databases.

Database groups let you organize access to databases based on whatever criteria you want. For example, based on a database's geographic location, its ownership, or its purpose. Within each group are assigned registered databases (or databases). Also within each group are assigned cockpit users. When these users log in to SAP HANA cockpit, they're able to monitor each of the databases within the group, as well as see aggregate data for the group.

You can assign the same registered database and cockpit user to multiple database groups. A cockpit user not assigned to at least one database group or assigned to a database group that has no registered databases assigned can log in to SAP HANA cockpit but have no access to any databases. A registered database unassigned to at least one database group can't be managed by any cockpit user.

3.3 SAP HANA Cockpit User Management

Manage SAP HANA cockpit user credentials and the database groups they can access.

SAP HANA Cockpit Users vs. Database Users

SAP HANA cockpit users and database users are two distinct users. Though they share similar concepts and terminology, such as roles and groups, they each have their own distinct meaning and function.

You use a cockpit user to log on to SAP HANA Cockpit or Cockpit Manager. They are created and managed in Cockpit Manager. Cockpit users are assigned cockpit roles, which dictate what tasks they can perform in SAP HANA cockpit or Cockpit Manager. A cockpit user that has been assigned the role to manage cockpit users can add and remove roles from any cockpit user except themselves. They can change the password of any user, including themselves. Cockpit users are assigned to database groups, which contain the registered databases that will be available in SAP HANA cockpit. Once logged in to cockpit, to access one of the registered databases, you must supply the credentials of a database user.

Each registered database has its own set of database users. A database user can monitor and manage a database in SAP HANA cockpit. Database users are managed in each registered database, using roles and privileges to control the tasks each can perform within the database. Just as you can't use cockpit credentials to access a registered database, you can't log on to cockpit or Cockpit Manager using database credentials.

SAP HANA Cockpit Master User

During cockpit installation, a master user, COCKPIT_ADMIN, is automatically created. Its password is the master password specified during the installation process. This master user is assigned all administrator roles, and can perform all tasks in the Cockpit Manager.

Cockpit roles were introduced as of SAP HANA 2.0 SPS 01. If your `COCKPIT_ADMIN` user was created during the installation of an earlier version, then you may wish to assign it additional roles if you want to continue to access all aspects of the cockpit and the Cockpit Manager. You must log out and then back on again to have these new roles take effect for the `COCKPIT_ADMIN`.

Business Users

Business users are users created outside of cockpit to access other applications, such as SAP HANA Client or XSA. If these users now need to access cockpit, rather than create a new user in cockpit, you can add the existing user to Cockpit Manager.

3.3.1 Create or Enable an SAP HANA Cockpit User

You can create new SAP HANA cockpit users, enable an existing cockpit user whose cockpit access has been revoked, or allow existing business users to access the SAP HANA cockpit.

Prerequisites

You are logged on to SAP HANA Cockpit Manger as a cockpit user that has been assigned the `Cockpit User Administrator Role`.

Context

A new cockpit user can't see any databases in SAP HANA cockpit until someone with the `Cockpit Database Administrator` role assigns them to at least one database group that contains at least one registered database. An existing cockpit user whose access to cockpit was revoked must be re-enabled before it's visible in the Cockpit Users list in Cockpit Manager.

Procedure

1. In SAP HANA Cockpit Manager, select *Cockpit Users*.
2. Select *Create User*.
3. Do one of the following:
   - Enter a user name, password, and e-mail address for the user.
   - Select the checkbox to enable an existing cockpit or business user to access the cockpit, then select the user from the dropdown list. The password and email-address for the user are that last set for the user in the outside application for a business user, or in Cockpit Manager for a cockpit user.
4. If the user it to use Kerberos single sign-on with SAP HANA Cockpit, select *SSO with Kerberos* and specify the external Kerberos credentials.
5. Assign one or more cockpit roles to the user. There's no single, all encompassing, administration role. Access to specific sections within Cockpit Manager requires specific roles. For example, to manage databases, groups, and users would require both the `Cockpit Database Administration Role` and the `Cockpit User Administration Role`. If a cockpit user is never expected to log in to SAP HANA cockpit, then there's no need to grant them the `Cockpit User Role`.

Role	Permits access to
Cockpit Administrator Role	The *Cockpit Settings* section of the Cockpit Manager, where they can configure cockpit settings.
Cockpit Database Administrator Role	The *Issues with Technical Users*, *Registered Databases*, and *Database Groups* sections of the Cockpit Manager, where they can

Role	Permits access to
	register databases, create database groups, and assign cockpit users to database groups.
Cockpit User Administrator Role	The *Cockpit Users* section of the Cockpit Manager, where they can create and manage cockpit users.
Cockpit User Role	The SAP HANA cockpit, where they can view and monitor all databases in any assigned database groups.
Cockpit User Role and allow this user to register databases	The SAP HANA cockpit and the *Registered Database* section of the Cockpit Manager.
Cockpit User Role and Cockpit Configuration Template Administrator Role	The SAP HANA cockpit and the *Configuration Templates* section on the SAP HANA cockpit *Home* page.
Cockpit Troubleshooting Role	The *XSA Logs* section of the Cockpit Manager. **i Note** To view XSA Logs, you also need to assign the XSA role of Space Auditor.

6. (Optional) Click and select one or more database groups to which this user is a member. You can add the user to database groups later if needed, but a cockpit user is unable to access any registered database until this designation is done.
7. Select *Create User* or *Enable User*, depending on your action from step 3.

Next Steps

- If you didn't assign the user to an existing database group, do it now.
- If the database group didn't already exist, create it now and add the user.

Related Information

Add or Remove Database Groups for Cockpit Users [page 42]
Creating a Database Group [page 62]
Configure Kerberos for SAP HANA Database Hosts
Configure SSO with SPNEGO and Kerberos for XS Advanced Applications
SAP Note 1813724
SAP Note 1837331

3.3.2 Add or Remove Database Groups for Cockpit Users

Add or remove the database groups that cockpit users are a member of.

Prerequisites

- You are logged on to SAP HANA Cockpit Manger as a cockpit user that has been assigned the `Cockpit User Administrator Role`.
- The database groups the cockpit user belongs to already exist. See *Create a Database Group*.

Context

In order for cockpit users to monitor and manage a registered database, you need to assign them to the database group to which the database belongs.

This task can also be done in the *Database Groups* section of the Cockpit Manager. See *Add or Remove Cockpit Users in Database Groups*.

Procedure

1. In SAP HANA Cockpit Manager, select *Cockpit Users*.
2. Select the user whose access you want to add or remove from a group.
3. If necessary, then click *Database Groups* to display the list of the database groups the user is already a member of.
4. To add a group, select *Grant Access to Database Groups*, select one or more groups, and then click *OK*.
5. To remove a user's membership in a group, click the ⊗ *(Delete)* icon beside the group.

Related Information

Adding or Removing Cockpit Users in Database Groups [page 65]
Creating a Database Group [page 62]

3.3.3 Editing Settings for a Cockpit User

You can modify some settings for an SAP HANA cockpit user.

Prerequisites

You are logged on to SAP HANA Cockpit Manger as a cockpit user that has been assigned the `Cockpit User Administrator Role`.

Context

Except for deactivating or activating a cockpit user, you can't edit multiple cockpit users at the same time.

Procedure

1. In SAP HANA Cockpit Manager, select *Cockpit Users*.
2. Select the user whose access you want to modify.
3. To see the roles assigned to the user, click *Roles*.
4. To change the password, single-sign-on credentials, or assigned roles for the user, click *Edit*, make any changes, and then click *Save*.

   You can't change the roles assigned to the active cockpit user.

Role	Permits access to
Cockpit Administrator Role	The *Cockpit Settings* section of the Cockpit Manager, where they can configure cockpit settings.
Cockpit Database Administrator Role	The *Issues with Technical Users*, *Registered Databases*, and *Database Groups* sections of the Cockpit Manager, where they can register databases, create database groups, and assign cockpit users to database groups.
Cockpit User Administrator Role	The *Cockpit Users* section of the Cockpit Manager, where they can create and manage cockpit users.
Cockpit User Role	The SAP HANA cockpit, where they can view and monitor all databases in any assigned database groups.
Cockpit User Role and allow this user to register databases	The SAP HANA cockpit and the *Registered Database* section of the Cockpit Manager.
Cockpit User Role and Cockpit Configuration Template Administrator Role	The SAP HANA cockpit and the *Configuration Templates* section on the SAP HANA cockpit *Home* page.
Cockpit Troubleshooting Role	The *XSA Logs* section of the Cockpit Manager.

Role	Permits access to
	Note To view XSA Logs, you also need to assign the XSA role of Space Auditor.

5. To see the database group assignments for the user, click *Database Groups*.
   a. To add a group, select *Grant Access to Database Groups*, select one or more groups, and then click *OK*.
   b. To remove a user's membership in a group, click the ⊗ *(Delete)* button beside the group.

Related Information

Deactivate or Activate a Cockpit User [page 44]
Delete a Cockpit User or Revoke Cockpit Access [page 45]
Configure Kerberos for SAP HANA Database Hosts
Configure SSO with SPNEGO and Kerberos for XS Advanced Applications
SAP Note 1813724
SAP Note 1837331

3.3.4 Deactivate or Activate a Cockpit User

Temporarily deactivate or activate one or more SAP HANA cockpit users.

Prerequisites

You are logged on to SAP HANA Cockpit Manger as a cockpit user that has been assigned the `Cockpit User Administrator Role`.

Context

Deactivating a cockpit user prevents the user from logging in to SAP HANA cockpit. Credentials and privileges are retained in the deactivated account and are reapplied when you activate the user. You can deactivate or activate multiple users at the same time. If the statuses of the selected users are all the same, the applicable *Activate* or *Deactivate* button becomes available, but if the statuses are mixed, both buttons become available. Once you select an action, a confirmation prompt appears indicating how many of the users meet the change criteria. If you proceed, users not meeting the criteria are ignored.

You cannot change the status of the currently logged on user or a user authenticating with single sign-on (SSO).

Procedure

1. In SAP HANA Cockpit Manager, select *Cockpit Users*.
2. Select users:

To	Action
Select only one user	Click the user in the list.
Select multiple users	Click the *(Multi Select)* icon at the top of the Cockpit Users list to allow multi selected user. Click a check box to add or remove each user from the selection
Select all users	Click *Select All*. To unselect all users, click the ⊗ *(Multi Select)* icon.

3. Click *Deactivate* or *Activate*.

Results

The status of each selected user changes to reflect the status change.

3.3.5 Delete a Cockpit User or Revoke Cockpit Access

You can delete cockpit users or just revoke their access to cockpit.

Prerequisites

You are logged on to SAP HANA Cockpit Manger as a cockpit user that has been assigned the `Cockpit User Administrator Role`.

Context

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

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

Once a user's access to the cockpit is revoked, they don't appear in the Cockpit Users list. However, if you try to create a new user with the same users name, a message appears warning that the user name already exists. To

completely delete a cockpit user whose cockpit access has been revoked, you must first re-create the user and then completely delete them.

Procedure

1. In SAP HANA Cockpit Manager, select *Cockpit Users*.
2. Select the user whose access you want to delete or revoke.
3. Select *Delete*.
4. In the *Confirm Request* dialog:

Choose...	In order to...
Remove Access Only	Allow the business user to continue to exist outside the cockpit, but without cockpit access.
Delete User	Delete the business user from the cockpit and from any additional applications.
Cancel	Don't delete the user nor the cockpit access.

Results

The cockpit user disappears from the list. If you only revoked access for the user, to restore access, see *Create or Enable a SAP HANA Cockpit User*.

Related Information

Security Considerations for SAP HANA Cockpit [page 24]
Create or Enable an SAP HANA Cockpit User [page 40]

3.4 Registered Databases

A cockpit user with the `Cockpit Database Administrator` role can register and manage databases.

3.4.1 The Technical User

The `technical user` is a dedicated database user that the cockpit uses to collect health data for monitoring the database (such as information on alerts and system performance).

The `technical user` must be unique for each database within the same instance. If you reuse the `technical user`, then registration fails with a message that the system can't authenticate the user credentials.

Create the `technical user` during database registration by specifying the credentials of an administrator on that database with the ability to create user accounts. You can also create a `technical user` by using the *User and Role Management* card on the database *Database Overview* page as long as you have the user creation privilege.

The `technical user` should be exempt from password expiration policies and should be a dedicated account. Human users should never log in using the `technical user` account. At a minimum, the `technical user` requires the CATALOG READ system privilege and the SELECT privilege on the _SYS_STATISTICS schema. Additional privileges may be required, as additional optional permissions are enabled to enhance data collection. See *Security Considerations of SAP HANA Cockpit*.

If changing the `technical user` of a registered database, then you must create the new user first by using *User and Role Management* in SAP HANA cockpit, and then edit the registered database in SAP HANA Cockpit Manager, specifying the new technical user`technical user` credentials. You can't create a `technical user` when editing a registered database. Once the new technical user is assigned, if you don't need the old `technical user`, then deactivate it by using *User and Role Management*.

Resources with an Invalid Technical User

If there are issues with your database's `technical user` credentials, then you're notified via a message strip on your database in the *Database Directory* or *Cockpit Manager*.

This message persists until you provide the correct credentials by clicking *Investigate* on the message strip, or by clicking on the *Issues With Technical Users* option on the *Cockpit Manager*.

If all `technical user`s are valid, then the *Issues With Technical Users* option no longer appears on the *Cockpit Manager*.

Related Information

Security Considerations for SAP HANA Cockpit [page 24]

Create a Database User [page 349]

3.4.2 Register a Database

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

Prerequisites

- You are logged on to SAP HANA Cockpit Manager as a cockpit user that has been assigned the `Cockpit Administrator Role` or the `Cockpit User Role` with the option to register databases.
- A `technical user` account with the required privileges already exists or you can provide the credentials of a database user with user creation privileges.
- If you plan to encrypt the SAP Control or database connection, in SAP HANA XS advanced, then you have:
  1. Manually imported the server root certificates
  2. Trusted the certificates by using the command syntax `xs trust-certificate <ALIAS> -c <CERT_FILE>`
  3. Exported the certificates to the cockpit by using the commands `xs restage cockpit-hdb-service` followed by `xs restart service cockpit-hdb-svc`

  See *Import a Certificate for Encrypted Communication*.
- If you plan to add the database to a group during the registration process, then the group already exists. See *Create a Database Group*.

Context

To make a database available to a cockpit user, the registered database must be a member of at least one database group of which the cockpit user is a member.

Procedure

1. Click *Registered Databases*.
2. In the section:
   a. Indicate whether the name of the is system-generated (default) or user-defined.
   b. Enter the host name (for on-premise) or the instance endpoint (for cloud) for the to register. Use a fully qualified host name if possible.

      You can register a whose statistics server isn't running or is unreachable, but some cockpit features aren't available for that until its statistics server is started or recovers.
   c. Specify an identifier for the — the instance number or a port number (of the SAP HANA nameserver).

d. (For Instance Number only) Select *Multiple containers* and then specify the type of database. If you selected Tenant database, then enter the tenant database name.
e. (Optional) Enter a description of the .

3. In the Connection section, choose whether to encrypt the cockpit's connections to SAP Control (for starting and stopping) and to the database. For information about obtaining certificates, see *Certificate Management in SAP HANA* in the *SAP HANA Security Guide*.
   - If you encrypt the SAP Control connection, then you're allowing a secure connection (HTTPS) to SAP Control (provided that you've met the prerequisite of importing the trusted certificates to the cockpit).
   - If you encrypt the database connection using a secure JDBC connection, then choose whether to validate the certificate. This option lets you stipulate whether to verify that the remote server is trusted by the cockpit. Deselect the checkbox if the SAP HANA database has a certificate that differs from the one currently imported, or if you haven't imported the certificate from the SAP HANA database into XS advanced. However, the recommendation is that you instead import a certificate for encrypted connections.
   - Optionally, you can enter a hostname to override the one in the certificate. Entering a hostname avoids the validation failure that could result from the certificate's hostname differing from the hostname that cockpit uses to connect. For example, if there's a host alias, or a short hostname instead of a fully qualified domain name.
4. In the `technical user` section:

Option	Description
If the `technical user` already exists	Enter its user name and password. Cockpit Manager doesn't verify the specified technical user at this point. That happens later in the registration process. **i Note** The `technical user` must be unique for each within an instance. For example, if you've already registered the SYSTEMDB and are now registering the tenant , then the `technical user` for the tenant must be different from the SYSTEMDB.
If the technical user doesn't already exist	Click *Create New Technical User* and specify the following: 1. The user name and password of a administrator user with user creation privileges. 2. The user name and password for the new technical user. 3. (Optional) Select *Grant EarlyWatch Alert permission*. 4. Click *Create*. You get a message regarding the success or failure of the user creation. If creation failed, then the message indicates why.

5. (Optional) In the Groups section, select the groups that this belongs to. You can add a to a group after registration is complete.
6. (Optional) In the Contact section, enter contact information for the user responsible for the .
7. Select *Review*.
8. Go over the details on the Register Review page and use *Edit* to make changes.
9. When you're satisfied with the information on the , select *Register*.

a. If you specified the port number for the identifier, then specify the instance number, select *Multiple containers*, and then specify the type of database. If you selected Tenant database, then enter the tenant database name.
b. Select *Register.*

Results

- If registration is successful, then the newly registered appears in the registered list, and displays the software version. In SAP HANA Cockpit Manager, select
- If you're prompted to register the system using the SAP HANA Agent, then click *Yes*. Enter the instance number, select *Multiple containers*, and then specify the type of database. If you selected Tenant database, then enter the tenant database name. Click *Register*.
- If registration is successful, but Cockpit Manger is unable to establish a connection to the , the newly registered appears on the Registered list with a software version of `0.00.000.00.0 (UNKNOWN)`. Before you can access the , resolve the connection issue. See *Troubleshooting Registration Issues*.
- If the registration request fails because the specified information is invalid, then select *OK* to acknowledge the error. Select *Edit* in the section with the error, fix it, and scroll to the bottom of the page. Select *Review* and try registering the again. See *Troubleshooting Registration Issues*.

Next Steps

- If Cockpit Manger is unable to connect to the , but is able to register the , then warning messages regarding failed statuses and privileges relating to the technical user appear. These messages appear because the system is unable to validate the related information. Acknowledge the messages. The newly registered appears in the registered list, but the software version appears as `0.00.000.00.0 (UNKNOWN)`. Before you can use the , you must resolve the connection issues. See *Troubleshooting Registration Issues*.
- Verify that there are no issues with the technical user. See *The Technical User*.
- If you didn't assign the to a group during registration because the group didn't already exist, then create it now and add the new registered to the group.
- If the group exists but you didn't assign the to a group during registration, then assign it now.

Related Information

Troubleshoot Registration Issues [page 51]
Creating a Database Group [page 62]
Importing a Certificate for Encrypted Communication [page 32]
Adding or Removing Databases in a Database Group [page 64]
Deactivate a Database User [page 362]
Delete a Database User [page 363]
The Technical User [page 47]

3.4.2.1 Troubleshoot Registration Issues

database registration fails when Cockpit Manager can't establish a connection to a database.

Cockpit Manager can't always determine if the failure is due to an unreachable database (for example, the database is stopped) or that the specified database doesn't exist. Different messages appear as Cockpit Manger tries to complete the registration process, depending on the underlying issue.

If Cockpit Manger is unable to register the database with the information you provided, an error message appears. Registration stops, but you remain on the *Register Database Review* page, and your registration values remain available. You can edit and correct your entries and try registering the database again. Check the following.

Invalid host name entered.	If you entered an invalid host name, select *OK* to acknowledge the error. Select *Edit* in the database section, fix the host name, and scroll to the bottom of the page. Select *Review* and try registering the database again.
Unable to authenticate user credentials.	The `technical user` must be unique for each database within the same instance. When you try to reuse the `technical user`, registration fails with a message that the system can't authenticate the user credentials. Select *Edit* in the Technical User section. Enter a different database user name (with sufficient privileges). If there isn't an available `technical user`, click *Create new Technical User* and define a new user. Scroll to the bottom of the page, select *Review*, and try registering the database again.

If Cockpit Manger is able to register the database, but is unable to connect to it, the newly registered database appears on the Registered database list with a software version of `0.00.000.00.0(UNKNOWN)`. To troubleshoot the connection issue, check the following. Once a connection is established, the software version number is updated for the database.

The database being registered is stopped.	The database must be running to properly register it. Verify that the database is running and start it if necessary. Once Cockpit Manager can connect to the running database, the software version appears on the Registered database list.
The tenant name, system ID, or instance number is incorrect.	Review the database details and verify the information. If any of these values are incorrect, they can't be changed after registration. You must unregister the database and try registering it again.

If you can't identify and resolve the cause of the connection issue, try unregistering the database and then try registering again.

Related Information

Create a Database User [page 349]

3.4.3 Edit Database Settings

Once a database has been registered, you can modify some settings, including `technical user`, encryption, and single sign-on (SSO).

Prerequisites

You are logged on to SAP HANA Cockpit Manager as a cockpit user that has been assigned the `Cockpit Administrator Role` or the `Cockpit User Role` with the option to register databases.

Context

When you log in to Cockpit Manager, if the credentials for this database's `technical user` (the logon the cockpit uses to collect system health and version information) need to be updated, then an indicator in the *Database Directory* alerts you. This behavior happens because the `technical user`'s password has changed, or the `technical user` has been deleted. You can't monitor the database until you fix the technical user's credentials.

Procedure

1. Click *Registered Databases*.
2. Select the database you want to modify.
3. On the *Database Details* tab, click *Edit*.
4. Modify the settings as needed.
    - Technical user
        - To change the technical user, you must supply a database user with user creation privileges.
        - You can:
            - Enter a new technical user as long as it already exists in the database. The new technical user must already exist; you can't create a new user while editing a registered database. Each database within the same instance must be assigned a unique technical user. If you try to assign the same technical user to multiple databases, then you get a `Request Failed` message.
            - Grant and remove access to other services used by the technical user. Granting access automatically grants any additional privileges required for the feature. However, removing access doesn't remove the granted privileges from the user.
            - Change the password for the technical user. The technical user passwords entered for the database in cockpit manager and in cockpit are stored in different locations, and must always match. Changing the password in one location doesn't automatically update the other location.
    - Connection

- If you encrypt the SAP Control connection, then you're allowing a secure connection (HTTPS) to SAP Control (provided that you've met the prerequisite of importing the trusted certificates to the cockpit).
- If you encrypt the database connection using a secure JDBC connection, then choose whether to validate the certificate. This option lets you stipulate whether to check that the remote server is trusted by the cockpit. You can use this option if you haven't imported the certificate from the SAP HANA database into XS advanced, or if the SAP HANA database has a certificate that differs from the one currently imported.
- Optionally, you can enter a hostname to override the one in the certificate. Providing a hostname avoids the validation failure that results from the hostname in a certificate differing from the hostname that cockpit uses to connect. For example, if there's a host alias, or a short hostname instead of a fully qualified domain name.
  For information about obtaining certificates, see *Certificate Management in SAP HANA* in the *SAP HANA Security Guide*.

- Single Sign-on (SSO)
  - To enable SSO, supply a database user with the TRUST ADMIN, CERTIFICATE ADMIN, and USER ADMIN privileges. Before enabling SSO, consider migrating the Personal Security Environment (PSE) file to an in-database store. When SSO is enabled, a new PSE file may be created, which may prevent cockpit access to stored certificates. See SAP Note 265666.
  - To enforce SSO, you must supply a database user with the TRUST ADMIN, CERTIFICATE ADMIN, and USER ADMIN privileges. When enforced, cockpit users must use SSO to access the database in SAP HANA cockpit. Otherwise, on the *Database Directory* page, cockpit users can choose whether to access this database with SSO or enter alternate database user credentials. Enforce SSO only after you have configured SSO on the database. See *Enforce Single Sign-On on the Database*.
- Data Collection (Status and Alert Counts, and Feature Data, EarlyWatch Alert Transmission and Collection)
  - You can modify the collection settings for a specific database by editing the details of that database. Doing so overrides the global settings for that particular database. Only perform changes if necessary.
  - To revert to the global setting, in database edit mode, select *Global Settings* (instead of *Database Override*).

5. When done, click *Save*.

Related Information

User Authentication and Single-Sign On
Enforce Single Sign-On on the Database [page 54]
SAP Note 265666
Override Data Collection for a Database [page 55]
SAP HANA Security Guide for SAP HANA Platform

3.4.3.1 Enforce Single Sign-On on the Database

Before you can enforce SSO through cockpit only, you must configure SSO on the database.

Prerequisites

- The cockpit users who access this database have been created.

Context

You can enforce single sign-on (SSO) user authentication on database running SAP HANA 2.0 SPS 01 or later.

Procedure

1. As a cockpit user who is a member of the registered database's database group. log in to SAP HANA cockpit.
2. Click the *Database Directory* tile.
3. Connect to the database as an existing database user with the TRUST ADMIN, USER ADMIN, and CERTIFICATE ADMIN privileges.
4. On the *Database Overview* page, on the *User & Role Management* card, select *Manage Users*.
5. Select a user to use SSO and click *Edit*.
6. Set the JWT mappings.
7. Return to the *Database Directory* page and connect to the database as the newly configured database user.
8. Return to the Cockpit Manager, edit the database and enforce SSO.

Next Steps

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

3.4.3.2 Override Data Collection for a Database

You can modify the collection settings for a specific database by editing the details of that database. Doing so overrides the global settings for that particular database, so only perform changes if necessary.

Prerequisites

You know the password of a cockpit user that has been assigned the `Cockpit Database Administrator` role.

Procedure

1. On the Cockpit Manager page, click *Registered Databases*.
   View all the systems known to the SAP HANA cockpit on the *Registered Databases* sidebar.
2. In the left pane, select the database whose settings you want to modify.
3. In the right pane, select the *Data Collection Settings* tab, and view the globally set data collection settings.
4. Click *Edit*.
5. In either the *Status and Alert Counts* or the *Database and Feature Data* sections, or both, select *Database Override*.
6. Opt to disable the data collection or to modify the settings.
7. Select *Save*.

> **i Note**
>
> To revert to the global setting, in database edit mode, toggle to *Global Settings* (instead of *Database Override*).

Related Information

Setting Data Collection [page 80]

3.4.4 Unregister a Database

Remove one or more database from SAP HANA cockpit.

Prerequisites

You are logged on to SAP HANA Cockpit Manager as a cockpit user that has been assigned the `Cockpit Administrator Role` or the `Cockpit User Role` with the option to register databases.

Context

Removing a database also removes it from any associated database groups.

Procedure

1. Click *Registered Databases*.
2. Select the database to remove and click *Unregister*.

Related Information

Register a Database [page 48]

3.4.5 Export Databases

Export registration information about databases to a JSON file, which you can then import into other systems.

Prerequisites

- You are logged on to SAP HANA Cockpit Manager as a cockpit user that has been assigned the `Cockpit Administrator Role` or the `Cockpit User Role` with the option to register databases.
- The databases you're exporting are running and available on the network.

Context

You can export multiple databases at the same time. Optionally, you can include the `technical user` and contact information in the export file. When you include the technical user, the password is not included.

Procedure

1. Click *Registered Databases*.
2. Click the ... *(More)* button, then click *Export Databases*.
3. Select the databases you want to export.
4. (Optional) Click to *Save technical user login names* in the export file.
5. (Optional) Click to *Save contact settings* in the export file.
6. Click *Save export file*.

Next Steps

By default, the export file is saved to the Downloads folder and named `ResourceList.json`. Copy the export file to a location accessible by the importing system, then use the file to import the databases to the new system. If you included the `technical user` in the file, you can:

- Edit the export file and add the corresponding password before initiating the import. The `technical user` and password values then automatically populate the fields during the import. See *Add the Technical User Password to the Export Json File*.
- Manually enter the `technical user` and password during the import.

Related Information

Import Databases [page 58]
Add the Technical User Password to the Export JSON File [page 60]

3.4.6 Import Databases

Add databases exported from other systems to the current system.

Prerequisites

- You are logged on to SAP HANA Cockpit Manager as a cockpit user that has been assigned the `Cockpit Administrator Role` or the `Cockpit User Role` with the option to register databases.
- You've created a `.json` or `.xml` export file on another SAP HANA system that specifies the databases to be imported, and you've copied that file to a location accessible to the importing system.
- If you plan to encrypt the SAP Control or database connection, in SAP HANA XS advanced, then you have:
  1. Manually imported the server root certificates
  2. Trusted the certificates using the command syntax `xs trust-certificate <ALIAS> -c <CERT_FILE>`
  3. Exported the certificates to the cockpit using the commands `xs restage cockpit-hdb-service` followed by `xs restart service cockpit-hdb-svc`.

  See *Import a Certificate for Encrypted Communication*.
- The databases you're importing are running and available on the network.

Context

The export file may contain the `technical user` and or the password for each database. See *Add the Technical User Password to the Export JSON File*. If included, the information automatically populates the fields on the *Authenticate* page. You can override the information, if needed, but if specified, the user name and password must exist in the database.

The specification of a `technical user` and password during import is optional, so you can leave the fields empty, but if the `technical user` is specified , you must either enter a password or clear the `technical user` before you can proceed with the import.

If you don't specify a `technical user`, once the database is registered, you will see warning messages in SAP HANA Cockpit Manger and SAP HANA Cockpit that the `technical user` is missing and you are prompted to add it.

Validation of the `technical user` information doesn't occur until the actual import begins. If the `technical user` is invalid, an error message appears. You can edit the import details and try the import again.

Procedure

1. Click *Registered Databases*.
2. Click the ••• *(More)* icon, then click *Import Databases*.

3. On the Import page:
   a. Choose the file type to be imported.
   b. Click *Browse* and choose the file.
   c. Specify whether to include collection settings.

      The *Include collection settings* option appears for export files created from a cockpit version SP 10 or later. It includes settings for statuses, alerts, databases, and so on, which could override your global cockpit settings. Don't select this option if you don't want to override your global cockpit settings.
4. Select the databases you want to register.
5. On the *Authenticate* page,

Option	Action
To use the displayed technical user and password information from the export file	Go to step 6 [page 59].
To specify a technical user and password for a database	Manually enter the value. You can override any value displayed, but the credentials must be valid in the database being imported.
To assign the same technical user and password to each database	Click *Use the same credentials for all selected database*. Enter a single technical user and password when prompted. The credentials must be valid in every database being imported.
To clear credentials and import all selected databases without a technical user defined	Click *Clear credentials and register without providing a technical user for all selected databases.*

   When importing multiple databases, you can use different options for the databases. For example, if you are importing three databases, you could leave the displayed user and password for the first database, override the values from the import file for the second database, and clear the technical user and password completely for the third database.

6. Click the *Step 4* button.

   The *Step 4* button does not appear until both the technical user and password fields for each database either have values defined or are empty.
7. (Optional) If you included contact information in the export file, the values automatically appear. You can override any value displayed.
8. (Optional) Select a database group. All of the selected databases will be assigned to the same group.
9. Click *Review*. Click *Edit* to make any changes.
10. Click *Import Databases*.

    A message appears to indicate the status of each database's import. If the import failed for a database, you can edit the values and try again.

Next Steps

- (Optional) If you didn't specify a database group during the import, then specify one now.
- If you imported databases without specifying a technical user, the health data for monitoring databases will not be collected until a valid technical user is specified for the database. Add the technical user now.

Related Information

The Technical User [page 47]
Export Databases [page 56]
Importing a Certificate for Encrypted Communication [page 32]
Add the Technical User Password to the Export JSON File [page 60]
Adding or Removing Databases in a Database Group [page 64]

3.4.6.1 Add the Technical User Password to the Export JSON File

You can add the technical user password to the export file before importing the databases.

Context

When exporting a database, you had the option to include the `technical user` in the file. When included, the corresponding password is not included in the file, but the entry *technicalPwd* followed by an empty string is add as a placeholder for the password value.

You can edit the export file and add the password for the technical user fro any database before starting the import. The `technical user` and password must existing in the databases. If there are multiple databases in the export file, you don't need to specify a password for every technical user.

> **i** Note
>
> The password will be visible in the file, so it is recommended that you keep the file in a secure location, deleting the file once the import is complete.

Procedure

1. Open the export JSON file in a text editor.
2. Starting at the beginning of the file, search for *technicalPwd":*.
3. Add the `technical user` password to the empty string immediately following the search string. For example: *"technicalPwd":"password1234"*.
4. Repeat steps 2 [page 60] and 3 [page 60] to add additional passwords.
5. Save and close the file.

Next Steps

Perform an import.

Related Information

Import Databases [page 58]

3.5 Database Groups

You can create, populate, or remove the groups used to grant cockpit users access to registered databases.

A database group—a named set of one or more registered databases—controls management and monitoring privileges. You assign cockpit users and registered databases to a database group, enabling the user to monitor and manage the group's databases through the cockpit. A cockpit user can only access registered databases that are assigned to the same database group.

Database groups let you:

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

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

Each database also belongs to an auto-generated group. These auto-generated groups of databases (Production, Test, Development) are based on the system usage type of each database. System usage type is configured during system installation, or later using the `global.ini` file with the `usage` parameter in the `system_information` section.

You can hide the auto-generated groups if you don't want to use them. See *Displaying Auto-Generated Database Groups*.

You can't assign cockpit users to an auto-generated group, or to individual databases.

Related Information

Displaying Auto-Generated Database Groups [page 83]

3.5.1 Creating a Database Group

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

Prerequisites

- You are logged on to SAP HANA Cockpit Manager as a cockpit user that has been assigned the `Cockpit Database Administrator Role`.

Context

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

Procedure

1. In SAP HANA Cockpit Manager, select *database Groups*.
2. Click *Create Group*.
3. Enter a name for the group.

   The name can contain uppercase and lowercase letters, the digits 0 through 9, underscores, hyphens, periods, and spaces and must be unique.
4. (Optional) Enter a description for the group.
5. (Optional) Click  and select the registered databases to add to the new group. You can add registered databases to a group later if needed.
6. (Optional) Select *Register other databases after the group is created* so that as soon as the group is created, you can begin registering a new database. The new database is automatically added to the new group, but you can override this behavior if needed.
7. (Optional) Click  and select cockpit users who access the new group. You can add cockpit users to a group later if needed.
8. Click *Create Group*.

   If you selected *Register other databases after the group is created* during creation, then the group is created and then the *Register Database* page opens. Otherwise, the new group appears in the Database Groups list.

Next Steps

- If the registered databases or cockpit users didn't already exist when creating the database group, then create them now and add them to the new group.

- If you didn't assign existing users or registered databases to the new group during registration, then do it now.

Related Information

Register a Database [page 48]
Create or Enable an SAP HANA Cockpit User [page 40]
Adding or Removing Databases in a Database Group [page 64]
Adding or Removing Cockpit Users in Database Groups [page 65]

3.5.2 Deleting a Database Group

Remove a database group from the SAP HANA cockpit.

Prerequisites

You are logged on to SAP HANA Cockpit Manager as a cockpit user that has been assigned the `Cockpit Database Administrator Role`.

Context

Any cockpit users who had access to databases solely through the delete database group will no longer able to access them.

Procedure

1. In SAP HANA Cockpit Manager, select *database Groups*.
2. Select the group you want to remove.
3. Click *Delete* and confirm the deletion.

3.5.3 Adding or Removing Databases in a Database Group

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

Prerequisites

- You are logged on to SAP HANA Cockpit Manager as a cockpit user that has been assigned the `Cockpit Database Administrator Role`.
- The database to be added is already registered. See *Register a Database*.

Context

Auto-generated groups don't appear on the *Database Group* page and you can't modify or delete them. Register a database to add it to a group.

Procedure

1. In SAP HANA Cockpit Manager, select *database Groups*.
2. Select the group you want to add or remove a database from.
3. If necessary, then click *Database* to display the list of the group's databases.
4. To add a database, click *Add Database*, select one or more databases, and then click *OK*.
5. To remove a database, click the ⊗ *(Delete)* icon beside the database.

Related Information

Register a Database [page 48]

3.5.4 Adding or Removing Cockpit Users in Database Groups

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

Prerequisites

- You are logged on to SAP HANA Cockpit Manager as a cockpit user that has been assigned the `Cockpit Database Administrator Role`.
- The cockpit users to be added to the group already exist. See *Create or Enable a Cockpit User*.

Context

Auto-generated groups don't appear on the *Database Group* page and you can't modify or delete them.

Procedure

1. In SAP HANA Cockpit Manager, select *database Groups*.
2. Select the group you want to add or remove a database from.
3. If necessary, then click *Cockpit Users* to display the list of the group's users.
4. To add a user, click *Add User*, select one or more users, and then click *OK*.
5. To remove a user, click the ⊗ *(Delete)* icon beside the user.

Related Information

Create or Enable an SAP HANA Cockpit User [page 40]

3.6 Managing Databases, Users, and Groups with the Cockpit APIs

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

Using Cockpit APIs from an External Program

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

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

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

Obtaining a Service Key

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

Run the `xs create-service-key` command to generate a service key:

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

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

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

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

```
Getting service key "foo-cockpit-uaa" for service instance "cockpit-uaa" ...
{
"tenantmode" : "dedicated",
"clientid" : "sb-cockpit!i1",
"verificationkey" : "<REDACTED>",
"xsappname" : "cockpit!i1",
"identityzone" : "uaa",
"identityzoneid" : "uaa",
"clientsecret" : "<REDACTED>",
```

```
"url" : https://host:30032/uaa-security
}
OK
```

Change it to look similar to this:

```
{
  "tenantmode" : "dedicated",
  "clientid" : "sb-cockpit!i1",
  "verificationkey" : "<REDACTED>",
  "xsappname" : "cockpit!i1",
  "identityzone" : "uaa",
  "identityzoneid" : "uaa",
  "clientsecret" : "<REDACTED>",
  "url" : https://host:30032/uaa-security
}
```

Obtaining an OAuth Token

OAuth tokens are granted to cockpit users, so you must identify a cockpit user you can employ to run the cockpit APIs. To find the role needed for the API you want to run, see Using Cockpit POST APIs [page 68] or Using Cockpit GET APIs [page 71].

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

```
POST https://host:30032/uaa-security/oauth/token HTTP/1.1
Host: localhost:8080
Accept: application/json
Authorization: Basic YXBwOmFwcGNsaWVudHNlY3JldA==
"grant_type=password&username=marissa&password=koala&clientid=<clientid>
&clientsecret=<clientsecret>"
```

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

```
{
  "access_token":"2YotnFZFEjr1zCsicMWpAA",
  "token_type":"bearer",
  "expires_in":3600
}
```

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

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

Invoking Cockpit APIs

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

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

Related Information

Using Cockpit POST APIs [page 68]
Using Cockpit GET APIs [page 71]

3.6.1 Using Cockpit POST APIs

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

For instructions on setting up programmatic access to the cockpit APIs, see Managing Databases, Users, and Groups with the Cockpit APIs [page 66].

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

Success

If an API succeeds, it returns:

- Status 200
- A JSON response in this form:

  Sample Code

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

Failure

If an API fails, it returns:

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

Sample Code

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

SAP HANA Cockpit POST APIs

Each API listed in the POST APIs table accepts HTTP POST operations with arguments passed in the body in JSON format. Each API returns data in JSON format. These POST APIs require you to use an authentication token as described in Managing Databases, Users, and Groups with the Cockpit APIs [page 66]. Because out-of-the-box web browsers don't support POST operations, you need a suitable plug-in to use or test the cockpit's POST APIs with a browser.

POST APIs

API	What It Does	Input Parameters	Response[1]	Required Role

`/registration/SystemRegister`	Registers an SAP HANA resource with the cockpit using JDBC. Post to the API with parameters in the body in JSON format (`Content-Type=application/json; charset=UTF-8`).	`resourceNamehostName`
		`instanceNumber` or `port`
		`techUser`
		`techUserCredentials`
		`isMultiTenant` – can be omitted if you specify `port` instead of `instanceNumber`
		`databaseName` – can be omitted if you specify `port` instead of `instanceNumber`
		`security (encryptJDBC, validateServerCertificate, hostNameInCertificate)`
			`resid` – resource ID of newly registered resource	Cockpit Resource Administrator or Cockpit Power User

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

API	What It Does	Input Parameters	Response[1]	Required Role
`/registration/SystemUnregister`	Unregisters an SAP HANA resource. Post to the API with parameters in the body in JSON format (`Content-Type=application/json; charset=UTF-8`).	`resid` – integer, required. The unique resource ID previously returned by a registration API.		Cockpit Resource Administrator or Cockpit Power User
`/registration/ResourceUnregister`	Unregisters an SAP HANA resource. Deprecated—use `SystemUnregister` instead.	`resid`		Cockpit Resource Administrator
`/group/GroupCreate`	Creates a new resource group	`groupName` `groupDescription`	`groupId`	Cockpit Resource Administrator
`/group/GroupDelete`	Deletes a resource group	`groupId`		Cockpit Resource Administrator

`/user/CockpitUserCreate`	Creates a new cockpit user	`username`
		`password`
		`email`
		`roleCollections[]`
			`cockpitId` – ID of newly created user	Cockpit Administrator

`/user/CockpitUserDelete`	Deletes a cockpit user	`deleteFromUAA`
		`userId`
		`username`
				Cockpit Administrator

`/group/GroupResourceAdd`	Adds a resource to a group	`groupId` `resourceId`		Cockpit Resource Administrator
`/group/GroupResourceRemove`	Removes a resource from a group	`groupId` `resourceId`		Cockpit Resource Administrator
`/group/GroupUserAdd`	Adds a user to a group	`groupId` `userId`		Cockpit Resource Administrator

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

API	What It Does	Input Parameters	Response[1]	Required Role
`/group/GroupUserRemove`	Removes a user from a group	`groupId` `userId`		Cockpit Resource Administrator

3.6.2 Using Cockpit GET APIs

The SAP HANA cockpit APIs provide information about registered databases and database groups. The GET APIs don't create, delete, or change anything in the cockpit.

Prerequisites

- You must be authenticated as a SAP HANA cockpit user to invoke the APIs. Each API with "Get" in its name (RegisteredResourcesGet, for example) is protected with both authentication and authorization checks.
- You must have an OAuth token to invoke API calls against the cockpit-landscape-svc. To obtain a token, follow the instructions in Managing Databases, Users, and Groups with the Cockpit APIs [page 66].
- If you're using a browser to test a GET API, you must go through the app-router, using either the cockpit-admin-web-app or cockpit-web-app URL.

RegisteredResourcesGet

Returns information about the databases registered in SAP HANA cockpit.

Prerequisites

You must have the COCKPIT_RESOURCE_ADMIN role or the COCKPIT_POWER_USER role. The information returned depends on your role:

- If you have the COCKPIT_RESOURCE_ADMIN role, then RegisteredResourcesGet returns information on all databases registered with this cockpit.
- If you have the COCKPIT_POWER_USER role, then RegisteredResourcesGet returns information on all databases registered by you with this cockpit.

Returns

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

Syntax

You can invoke GET APIs in two ways:

- Through the cockpit-web-app port via the app-router. The call is redirected to the XSA sign-in page if it doesn't present an app-router cookie indicating authentication status.

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

This form of invocation through the app-router is ideal for testing with a web browser, but not ideal for programmatic calls.

- Against the cockpit-landscape-svc endpoint. Calls using this method must present a valid authentication token.

API Endpoints for RegisteredResourcesGet

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

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

RegisteredResourcesGet supports pagination and therefore returns a maximum of 100 rows for one page. If you have more than 100 registered databases, call `/cp/admin/RegisteredResourcesGet or /resource/RegisteredResourcesGet` to get the top 100 registered databases, then use the OData token `$top` and `$skip` to get the next 100 registered databases, for example:

```
/cp/admin/resource/RegisteredResourcesGet?$skip=100&$top=100
```

A successful response is a JSON object in this format:


```
{
    "result": [
        {
            "BuildNumber": "1522210459",
            "CertificateHostName": "",
            "CollectionConfigurations": [],
            "Connections": [
                {
                    "Host": "host.domain.com",
                    "IsSAPControlAuthenticated": false,
                    "PortType": "INSTANCE",
                    "PortValue": 1,
                    "Role": "MASTER",
                    "SAPControlUserName": null
                }
            ],
            "CreatedBy": "COCKPIT_ADMIN",
            "DatabaseName": "DBNAME",
            "Designation": "CUSTOM",
            "EncryptedJDBC": false,
            "EncryptedSAPControl": false,
            "GroupCount": 1,
            "HardwarePlatform": "30A8S20Q0B",
            "Host": "host.domain.com",
            "HostName": "host.domain.com",
            "OSVersion": "SUSE Linux Enterprise Server 12.1",
            "PatchLevel": 0,
            "Port": 1,
            "PortType": "INSTANCE",
            "ResIcon": "HANA_TENANTDB.gif",
            "ResKey": "RESKEY_HANA_MDB_TENANT",
            "ResValue": "HANA_MDB_TENANT",
            "ResourceDescription": "",
            "ResourceId": "ResourceId",
            "ResourceName": "DBNAME@SID",
            "ResourceOwnerDetail": "details",
            "ResourceOwnerEmail": "",
```

```
            "ResourceOwnerName": "",
            "ResourceUniqueId": "0f208532-8fcc-4aef-9bd1-40470aa03ad8",
            "SAPUser": null,
            "SSOEnabled": false,
            "SSOEnforced": false,
            "SSOSupported": true,
            "ServicePack": 30,
            "SystemName": "SID",
            "TechnicalUser": "SYSTEM",
            "ValidateServerCertificate": false,
            "Version": "2.00.030.00.1522210459 (hanaws, 2018.13.0)",
            "VersionMajor": 2,
            "VersionMinor": 0
        }
    ]
}
```


Each object in the array represents a registered database. The object returned has the following properties:

Property	Description
ResourceId	The internal ID the cockpit uses to identify this database. Other cockpit APIs require this ID as a parameter.
ResourceUniqueId	The unique internal identifier of the database. This value is either specified at registration time or read from the system being registered where possible. If not supplied by the caller or the database, a random value is generated but it isn't guaranteed to be unique.
ResourceName	The name of the database. For SAP HANA systems, the default name takes the form <DB>@<SID> – for example, DB1@HA0.
ResourceDescription	The description of the database optionally provided at registration time.
CreatedBy	The cockpit user who registered the database. (CreatedBy isn't the same as the database owner, following).The description of the database optionally provided by the registering user at registration time.
SystemName	The system name. For SAP HANA systems, it's the SID (HA0, for example).
DatabaseName	The database name, for example, DB1.
Designation	The designation or usage type of the database. For SAP HANA systems it can be PRODUCTION, DEVELOPMENT, TESTING, or CUSTOM.
Version	The database's full version string.
VersionMajor	The major version of the database as an integer (for example, 2)
VersionMinor	The minor version of the database as an integer (for example, 0)
ServicePack	The service pack of the database as an integer (for example, 20)
PatchLevel	The patch level of the database as an integer.
BuildNumber	The build number of the database.
ResKey	The description of the database optionally provided by the database key for the type of the database (for example "RESKEY_HANA_MDB_SYSTEM").
ResValue	The type of the database (for example, HANA_SYSTEM).
ResIcon	A path to an icon used to render the database type (not currently used)

Property	Description
OSVersion	The version of the operating system where the database is running
HardwarePlatform	The platform where the database is running (for example, VMware Virtual Platform).
ResourceOwnerName	Optional owner name. The owner isn't necessarily the same person who registered the system and doesn't have to be a cockpit user. This person the contact for questions or problems.
ResourceOwnerEmail	Optional. The e-mail address of the owner.
ResourceOwnerDetail	Optional. Additional details on the database owner. Details could include work hours, location, or phone number.
SSOSupported	If this value is true, this database supports single sign-on from the cockpit.
SSOEnabled	If this value is true, single sign-on from the cockpit is turned on.
SSOEnforced	If this value is true, single sign-on is the only available authentication method from the cockpit.
EncryptedJDBC	If this value is true, communication with this database uses an encrypted database connection.
EncryptedSAPControl	If this value is true, communication with the host agent managing this database uses TLS.
ValidateServerCertificate	If this value and EncryptedJDBC are both true, the client validates the server (database) using pre-installed certificates.
CertificateHostName	If set, the connection uses this value instead of the host name provided in the server certificate for encrypted database connections
TechnicalUser	The name of the `technical user` the cockpit uses for this database.
Host (& HostName)	The main host name of the database. Duplicated in both fields for compatibility.
PortType	The type of the main connection to the database. Set to either "INSTANCE" or "SQL". If INSTANCE the port is an instance number. If SQL the port is a SQL port.
Port	The instance number or SQL port (for example, the indexserver SQL port) of the main connection to the database.
GroupCount	The number of groups this database belongs to.

Property	Description
Connections	All possible connections to the database (for scale-out systems or those systems with host aliases, for example). The main connection as dictated by Role is copied into the previous properties. An array of connection objects in this form:

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

GroupsForUserGet

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

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

Syntax

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

API Endpoints for GroupsForUserGet

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

A successful response is a JSON object in this format:

```
{
```

```
result: [
{
Type: 1,
Name: "GROUPNAME",
Description: "Description",
RunWithAlert: 1,
NotRunning: 0,
ResourceCount: 5
}
]
}
```

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

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

GroupResourcesGet

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

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

Syntax You can invoke GroupResourcesGet in two ways:

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

API Endpoints for GroupResourcesGet

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

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

GroupResourcesGet supports pagination and therefore returns a maximum of 100 rows for one page. If you have more than 100 registered databases included in the specified group that is visible to the current user, call `/cp/ls/group/GroupResourcesGet` or `/group/GroupResourcesGet` (with mandatory query parameters as either `?groupId=<id>` or `?groupDesignation=<string>`) to get the top 100 registered databases. Then, use the OData tokens `$top` and `$skip` to get next top 100 registered databases, for example:

`/cp/ls/group/GroupResourcesGet?$skip=100&$top=100` (with mandatory query parameters as either `&groupId=<id>` or `&groupDesignation=<string>`)

Parameters

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

GroupResourcesGet requires one of these query parameters:

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

A successful response is a JSON object in this format:

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

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


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

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

Property	Description
XSASupported	If this value is true, this database has a running XS advanced server.
SSOEnabled	If this value is true, single sign-on from the cockpit is turned on.
SSOEnforced	If this value is true, single sign-on is the only available authentication method from the cockpit.
IsAuthenticated	If this value is true, the current user is authenticated with this database.
IsAuthenticatedWithSSO	If this value is true, the current user is authenticated with this database using single sign-on.
RemoteUserName	The database user currently used to authenticate with the database.
Host (& HostName)	The main host name of the database. Duplicated in both fields for compatibility.
State	The state of the database (RUNNING or STOPPED, for example).
Availability	A score describing the availability of the database.
Performance	A score describing the performance of the database.
Capacity	A score describing the capacity of the database.
AlertCountHigh	The number of high-level alerts currently active.
AlertCountMedium	The number of medium-level alerts currently active.
PortType	The type of the main connection to the database. Set to either "INSTANCE" or "SQL". If INSTANCE the port is an instance number. If SQL the port is a SQL port.
Port	The instance number or SQL port (for example, the indexserver SQL port) of the main connection to the database.
GroupCount	The number of groups this database belongs to.

Property	Description
Connections	All possible connections to the database (for scale-out systems or those systems with host aliases, for example). The main connection as dictated by Role is copied into the previous properties. An array of connection objects in this form:

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

3.7 Cockpit Settings

In the Cockpit Manager, as a cockpit administrator, you can select *Settings* to configure data collection, proxy server settings, and the connection timeout period, and to control whether or not SAP HANA Cockpit displays auto-created groups.

3.7.1 Setting Data Collection

You can reconfigure the default, SAP HANA cockpit global settings for collecting monitoring data. This data includes database status, alert counts, and other data from registered databases.

Prerequisites

You're logged on to SAP HANA Cockpit Manger as a cockpit user that has been assigned the `Cockpit Administrator Role`.

Context

Changes to SAP HANA cockpit settings are global and apply to the SAP HANA cockpit it manages.

You can change the defaults to specify:

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

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

> **i** Note
>
> You can also modify the collection settings for a specific database by editing the details of that database. Doing so overrides the global settings for that particular database. See *Override Data Collection for a Database*.

> **→** Tip
>
> The cockpit can support 1000s of registered databases. If the Multi-Host Health Monitor displays 'Not Collected' for specific databases, you can investigate the collection service log for rejected collections and reconfigure the worker threads accordingly.

Procedure

1. In SAP HANA Cockpit Manager, select *Cockpit Settings*.
2. Select *Data Collection* and then *Edit*.
3. Make the necessary changes and then save them.

Related Information

Using XS CLI Commands to Troubleshoot the Cockpit [page 89]

3.7.2 Set Up a Proxy Server

As a cockpit administrator, you can optionally set up a proxy server to use with SAP HANA cockpit.

Prerequisites

You're logged on to SAP HANA Cockpit Manger as a cockpit user that has been assigned the `Cockpit Administrator Role`.

Context

Changes to SAP HANA cockpit settings are global and apply to the SAP HANA cockpit it manages.

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

For an HTTP(S) proxy, you can also specify exceptions that don't use the proxy host. Use the *No Proxy Host* field to enter the exceptions (addresses beginning with the strings you enter, separated by semicolons).

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

Procedure

1. In SAP HANA Cockpit Manager, select *Cockpit Settings*.
2. Select *Proxy* and then *Edit*.
3. Make the necessary changes and then save them.

3.7.3 Setting Connection Timeout

You can specify the length of time that the SAP HANA cockpit waits for a connection before initiating a timeout.

Prerequisites

You're logged on to SAP HANA Cockpit Manger as a cockpit user that has been assigned the `Cockpit Administrator Role`.

Context

Changes to SAP HANA cockpit settings are global and apply to the SAP HANA cockpit it manages.

If a server connection is unresponsive, you may want to ensure that the cockpit doesn't wait for a response indefinitely. You can configure the following timeout periods:

Timeout period	Default
Standard database connect timeout	30 seconds
Long running tasks database connect timeout	48 hours
SAP Control connect timeout	15 seconds
SAP Control read timeout	30 minutes

Procedure

1. In SAP HANA Cockpit Manager, select *Cockpit Settings*.
2. Select *Connections* and then *Edit*.
3. Make the necessary changes and then save them.

3.7.4 Displaying Auto-Generated Database Groups

You can choose whether or not SAP HANA cockpit displays databases as part of auto-created database groups.

Prerequisites

You're logged on to SAP HANA Cockpit Manger as a cockpit user that has been assigned the `Cockpit Administrator Role`.

Context

Changes to SAP HANA cockpit settings are global and apply to the SAP HANA cockpit it manages.

Auto-generated groups contain databases that are based on the system usage type of each database (Production, Test, Development). System usage type is configured during system installation, or later using the `global.ini` file with the `usage` parameter in the `system_information` section.

You can choose to hide one or more of the auto-created groups. Opting to hide the auto-created groups doesn't affect the system usage type associated with the database. It simply prevents the cockpit from organizing the display of databases by auto-created group.

Procedure

1. In SAP HANA Cockpit Manager, select *Cockpit Settings*.
2. Select *Display* and then *Edit*.
3. Clear the check mark from the groups to hide.
4. Save your changes.

Related Information

Personalizing the Home Page [page 19]
Database Groups [page 61]

3.7.5 Enabling SSO with Kerberos

Enable single-sign-on (SSO) for Kerberos for SAP HANA Cockpit.

Prerequisites

- An SAP HANA user with the INIFILE ADMIN system privilege exists in SYSTEMDB.
- Your SAP HANA database is configured to use SSO with Kerberos.

Context

Once SSO with Kerberos is enabled, an SSO option appears on the User Details page for each Cockpit user.

Procedure

1. In SAP HANA Cockpit Manager, select *Cockpit Settings*.
2. Select *SSO with Kerberos* and then *Edit*.

3. Enable SSO and then enter the credentials of a user in SYSTEMDB with the INIFILE ADMIN privilege.

Next Steps

- Restart the XSA server for the configuration change to take effect. See *Restart SX Advanced* in the *SAP WEB IDE for SAP HANA Installation Troubleshooting Guide*.
- Enable SSO with Kerberos for applicable SAP HANA Cockpit users and specify the external Kerberos credentials.

Related Information

Create or Enable an SAP HANA Cockpit User [page 40]
Editing Settings for a Cockpit User [page 43]
Configure Kerberos for SAP HANA Database Hosts
Configure SSO with SPNEGO and Kerberos for XS Advanced Applications
Restart XS Advanced
SAP Note 1813724
SAP Note 1837331

3.8 User Notifications and Session Monitoring

Use Active Sessions to send pop-up notifications to SAP HANA cockpit users in real time and to monitor browser sessions.

Related Information

Monitoring Active Browser Sessions [page 86]
Sending a Notification to Logged-In Users [page 87]

3.8.1 Monitoring Active Browser Sessions

Learn how many users are connected to SAP HANA cockpit, what response time they're experiencing, and which part of the cockpit they're looking at.

Prerequisites

Your cockpit user has one of these roles:

- Cockpit Administrator
- Cockpit User Administrator

Context

Each active browser session represents a user. The Cockpit Manager's main screen lists the number of browser sessions next to the Active Sessions option.

On the Active Sessions screen, you can see:

- Which users are logged in to the cockpit. (Users logged in to the Cockpit Manager aren't included.)
- Each user's recent response time (latency) measured in milliseconds. Delays become noticeable at 150–200 ms.
- Each user's location in the cockpit. For example, `#Shell-home` is the Home screen and `#resourcedirectory-show` is the *Database Directory*.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit administrator or cockpit user administrator.

   You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation. The URL takes this form:

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

2. In the Cockpit Manager, select *Active Sessions*.
3. (Optional.) Enter a user's name in the Search Sessions field to filter the list of users.

Related Information

Sending a Notification to Logged-In Users [page 87]

3.8.2 Sending a Notification to Logged-In Users

Send a pop-up alert to all users logged in to SAP HANA cockpit.

Prerequisites

Your cockpit user has one of these roles:

- Cockpit Administrator
- Cockpit User Administrator

Context

To alert users to upcoming service disruptions or similar events, you can use Active Sessions to send two types of user notifications:

- *Inform*. These messages appear in the browser window for five seconds. Users don't need to react to them.
- *Interrupt*. These messages remain in the browser window until the user clicks to dismiss them.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit administrator or cockpit user administrator.

   You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation. The URL takes this form:

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

2. In the Cockpit Manager, select *Active Sessions*.
3. To send an inform message to all logged-in users:
   a. Click *Inform* (upper right).
   b. Enter your message.
   c. Click *Send Message*.
4. To send an interrupt message to all logged-in users:
   a. Click *Interrupt* (upper right).
   b. Select a type: *Information* or *Warning*.

      An information message displays an info symbol. A warning message displays an orange warning symbol.
   c. Enter a title.
   d. Enter your message.
   e. Choose text to appear on the button users must click to dismiss your message.
   f. Click *Send Message*.

Related Information

Monitoring Active Browser Sessions [page 86]

3.9 View Logs to Troubleshoot the Cockpit

Find entries in the SAP HANA cockpit operational logs.

Prerequisites

Your cockpit user has these roles:

- Troubleshooting (a cockpit role)
- Space Auditor (an XSA role - if you don't have the Space Auditor role, you can add it as described in step 2 [page 88].)

Context

View log entries to troubleshoot problems like missing monitoring data on cockpit pages, database registration failure, and connection timeouts.

Procedure

1. Connect to the Cockpit Manager and sign in as a user with the Troubleshooting role.

   You can reach the Cockpit Manager by entering the Cockpit Manager URL created during cockpit installation. The URL takes this form:

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

2. In the Cockpit Manager, select *XSA Logs*.

   If you don't have Space Auditor, an XS advanced role, a notification about the role appears. You can add Space Auditor with this `xs` command:

   ```
   xs set-space-role COCKPIT_ADMIN HANACockpit SAP SpaceAuditor
   Adding role 'SpaceAuditor' to user COCKPIT_ADMIN in space "SAP" of org
   "HANACockpit" ...
   OK
   ```

3. In the log viewer, there are two required fields for selecting and filtering log entries: *Application* and *Lines*.
   a. From the *Application* list, select a service—for example, cockpit-landscape-svc.

b. From the *Lines* list:
   - Select *All* to display all the entries in the log.
   - Select *Last* to specify the number of entries from the end of the log to display. For example, if you select *Last* and enter 10, the cockpit displays the last 10 entries in the log.
   - Select *Recent* to see the 25 most recent entries.
   - Select *Time Interval* to specify a range of time from which to display log entries. For example, to see log entries for the last four hours, select *Time Interval*. In the *Since* field, select today's date and set the time to four hours ago. In the *Until* field, select today's date and current time.

4. The log viewer has three optional fields: *Cockpit Area*, *Type*, and *Source*.
   a. (Optional) From the *Cockpit Area* list, select a component. Your selection here reduces the number of options in the *Application* list.
   b. (Optional) Select one or more log types from the *Type* list. To remove a log type, click the cancel icon ⊗ next to its name.
   c. (Optional) Select one or more log sources from the *Source* list. To remove a log source, click the ⊗ *(Cancel)*icon next to its name.
5. When you're satisfied with the settings in the filtering fields, click *Go*.

   The cockpit displays any log entries that meet you specifications.
6. (Optional) You can use the tools in the Log Lines row to search, download, sort, and customize the display of the log entries.
   - Search: Enter a full or partial search term in the *Search* field to display only log entries that contain your term.
   - Download: Click the ⤓ *(Download)* icon to download a `.txt` file containing the displayed log entries.
   - Sort: Click the ↑↓ *(Sort)* icon to choose sorting criteria for the log entries.
   - Customize display: Click the ⚙ *Table Personalization* icon to remove and reorganize the columns of the log entry table.

3.10 Using XS CLI Commands to Troubleshoot the Cockpit

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

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

Viewing Logs of Various Services

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

If you encounter issues with...	View the application log for...
Database registration, database group management, cockpit user management, or other Cockpit Manager issues	• cockpit-admin-web-app • cockpit-admin-ui-svc • cockpit-persistence-svc
Displaying or retrieving data, or issues related to specific SAP HANA cockpit applications	• cockpit-web-app • cockpit-hdbui-svc • cockpit-hdb-svc • cockpit-landscape-svc • cockpit-persistence-svc
Collections	• cockpit-collection-svc • cockpit-hdb-svc • cockpit-persistence-svc

→ Tip

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

Viewing Application Status

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

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

Related Information

View Logs to Troubleshoot the Cockpit [page 88]
Setting Data Collection [page 80]

SAP HANA Cockpit User Management [page 39]
Registered Databases [page 47]
SAP HANA Developer Guide for XS Advanced Model

4 Landscape Monitoring and Administration

For database groups to which you have access, you can monitor aggregated information representing each of the group's individual databases.

A database is an SAP HANA system, identified by a host and instance number. You can see data from multiple databases simultaneously—you can check the overall health of systems located within a data center or across your enterprise. You can drill down into status indicators for more detailed information.

Additional Functionality

The tiles each launch additional functionality. If you're the cockpit administrator user or a cockpit database administrator user, *Manage Cockpit* gives you access to the *Cockpit Manager*. Developers and administrators can also visually browse database objects (tables and schemas) and execute SQL statements.

Related Information

Monitoring Your Databases in the Database Directory [page 92]
Configurations and Configuration Templates [page 98]

4.1 Monitoring Your Databases in the Database Directory

Starting with an aggregate view of your registered databases in SAP HANA cockpit allows you to quickly discover any databases in your environment that have issues.

About the Database Directory

Use the *Database Directory* to view version and health information for all the databases in your landscape or to find a specific database instance. For each database, you can drill down for more information. Through the *Database Directory*, you can also specify the database user credentials required to access an individual database, which is necessary unless single sign-on is in effect for that instance.

When you first access the cockpit, the landscape level page displays important high-level information about all the SAP HANA databases to which you have been granted access. A database is identified by a host and instance number, and can be a single- or multi-host system. If you don't see any databases when you open the

cockpit, either there are no databases registered in the cockpit, or your cockpit database administrator hasn't assigned databases to you. See *Setting Up Cockpit with the Cockpit Manager* and *Setup and Administration with the Cockpit Manager*.

You can use the cockpit to monitor and manage more than one database, each running version SAP HANA 1.0 SPS 12 or later. Any database running version SAP HANA 2.0 SPS 01 or later is set in multiple-container mode by default. The cockpit can also monitor single-container systems running earlier versions of SAP HANA. When you access the *Manage Services* application, the operations you have the option to perform depend on whether you are displaying a tenant or a system database.

Database Groups

Each database belongs to a usage type group (Production, Test, Custom, or Development) depending on how it has been configured. Databases can also belong to one or more groups created by the cockpit database administrator.

You can also pay attention to the groups of databases and decide whether you need to navigate into a group to investigate potential issues, or you can click drill down to an individual database to obtain more details by clicking tabs, tiles, links, and numbers.

> → Tip
>
> Unless your administrator has enabled single sign-on, you need to connect to the database with a database user that has the system privilege CATALOG READ and SELECT on _SYS_STATISTICS.

Monitoring and Managing Database Health

The *Database Directory* provides a quick way to view essential health statistics for all the databases in your landscape. Clicking on a health statistic also provides a gateway to the application for that health component - for example, CPU usage - where you can view more information and manage configurations, and so on.

Default Database Directory View

Column	Description
Status	The status of the database. Can be one of: **Error** The database cannot be accessed. **Stopped** The database isn't running. **No SQL access** The database is running but isn't responding to SQL queries. **License Expired** The database is running but its software license has expired. If there's a license problem, click the database name. The *Database Overview* offers a link to the license manager. **Transitioning** The database is starting or stopping. **Running with issues** The database is running but has one or more high-severity alerts such as missed backups or an SSL certificate that's expired or about to expire. **Invalid technical user** The database is running, but SAP HANA cockpit is unable to monitor it because the credentials for the user that the cockpit employs to collect health data are incorrect or expired. If the technical user is invalid, click the database name. The *Database Overview* offers a link to the Cockpit Manager, where you (or an administrator with Cockpit Manager privileges) can update the technical user credentials. **Running** The database is running. Clicking on the status prompts you for the database credentials and then takes you to the *Manage Services* page.
Usage Type	Indicates what the database is used for. Can be one of: *Production*, *Custom*, *Test*, and *Development*.
Database	Your database's name and the host it is located on. Clicking on the database name takes you to the *Database Overview*.
Alerts	High and medium alerts for your database. Clicking on an alert icon takes you to the *Alerts* page.
Memory	Shows values for *Current* and *Peak* memory; the *Memory Allocation Limit*, and the *Threshold* for high and medium memory usage alerts. Clicking the memory graph takes you to the *Performance Monitor* where KPIs for your database's memory usage are displayed.

Column	Description
CPU	Shows values for *Current* CPU usage, the *Threshold* for high CPU usage alerts, and the *Start Time* and *End Time* of the 12 hour time range that is shown. Clicking the CPU graph takes you to the *Performance Monitor* where KPIs for your database's CPU usage are displayed.
Disk	Shows the *Current* percentage of disk usage, the *Threshold* for high and medium disk usage alerts, and the *Severity Level*, which can be one of *normal*, *warning*, or *critical*. Clicking the disk graph takes you to the *Performance Monitor* where KPIs for your database's disk usage are displayed.
Expensive Statements	The number of expensive statements executed by the database. Clicking on the number takes you to the *Expensive Statements* tab on the *Monitor Statements* page.
Group	The number of database groups the database belongs to. Click the number to display a list of database groups..
Type/Version	Indicates whether the database is a tenant or system database and displays the SAP HANA version. Clicking the version brings you to the *System Information*.
Credentials	When you click the link in the *Credentials* column, you can enter your SAP HANA database user name and password (unless single sign-on is enforced for this database). The cockpit securely encrypts and stores separate database credentials for each cockpit user; the database user name and password you enter can't be used by other cockpit users. See *Connect to a Database using Database Credentials*. You must be a user with the SYSTEM privilege to manage credentials. To save credentials for SAP Control to use for starting and stopping the database and restoring features, click Manage Credentials and enter the name and password of the database's <sid> adm OS user. If you wish to clear the credentials, click *Manage Credentials* and select *Delete the stored credentials for this database*. The cockpit securely encrypts and stores the credentials. You must be a user with the SYSTEM privilege to manage credentials.
SAP Control Credentials	Enter SAP Control logon credentials.
Database Management/SQL Console	If the database is a system database, then you can click *Database Management* to manage its tenant databases. If the database is a tenant database, then click on *SQL Console* to open a SQL console in the SAP HANA database explorer.

You can enable the following additional filters by clicking ⚙ *Table Personalization*:

Additional Database Directory Columns

Column	Description
Name	The name of the database instance. If you have chosen to use a system-generated name, then the name is comprised of the tenant or system database name along with the system name. For example, if your tenant name is TestDB and your system name is TestSystem, then your instance name is TestDB@TestSystem. If you are registering a system database for TestSystem, then your instance name is SYSTEDB@TestSystem.
Host	The host your database is located on.
System ID	The name of your database system.
Description	Any description that you've added to the database when you register it.
Availability / Performance / Capacity	High and medium alerts according to their corresponding Key Performance Area (KPA).
XSA Version	The XSA version, if installed, and any additional information, such as detailed information about version errors.
System ID@Host	The name of your database system and the host it is located on.
OS Version	The operating system version.
OS Kernel Version	The operating system kernel version.

Related Information

Memory Usage [page 139]
Monitoring, Analyzing, and Improving Performance [page 513]
Alerts [page 126]
Database Administration [page 427]
Managing Tenants from an SAP HANA System Database [page 371]
Backup and Recovery [page 591]
Security and User Management View [page 197]
Using the Database Overview Page to Manage a Database [page 109]
Getting to the Database Overview Page [page 125]

4.1.1 Connect to a Database With SSO or SAP HANA Credentials

Provide the credentials necessary to connect to a specific database using SAP HANA cockpit.

Context

The cockpit database administrator user may have used the Cockpit Manager configuration tool to enable cockpit to use the database's single sign-on (SSO) user authentication for a particular database (running SAP HANA 2.0 SPS 01 or later). If not, each cockpit user needs to provide the SAP HANA user credentials for their database to connect directly to the database. Each cockpit user should connect with different database user credentials. You can also connect using the `<sid>`adm user if you want the database to be able to access the SAP Control process (which involves starting and stopping the database, and restoring features).

> → Tip
>
> Unless your administrator has enabled single sign-on, you must connect to the database with a database user that has the CATALOG READ system privilege and the SELECT privilege on _SYS_STATISTICS.

> → Tip
>
> To improve system security, and to more effectively audit system activity, it is strongly recommended to set up a named database user account for each SAP HANA cockpit user in each registered system.

Procedure

On the *Database Directory* page, in the row displaying the database you want to connect to, choose one of the following:

- (Available for system databases only) Select the link in the *SAP Control Credentials* column, and enter the name and password of the database's `<sid>`adm user.
- Select the link in the *Credentials* column. The wording on the link and in the subsequent dialog depends on how the cockpit database administrator user has configured this database in the Cockpit Manager:

Database Configuration	Credentials Column displays...	Action
SSO has been enforced	*SSO Enforced*	Connect to the database through SSO. You don't need to enter database credentials.
SSO has been allowed	*SSO Enabled*, or the most recently used database user name, with a link to *Choose Authentication*	In the *Choose Authentication* dialog, you can choose to connect to the database through SSO or enter a different database user name and password.

Database Configuration	Credentials Column displays...	Action
SSO hasn't been allowed and you haven't previously connected to the database	The *Enter Credentials* link	In the *Enter Credentials* dialog, enter a database user name and password. The cockpit securely stores and encrypts the credentials for next time.
SSO hasn't been allowed but you've previously connected to the database	The most recently used database user name, with a link to *Manage Credentials*	In the *Manage Credentials* dialog, you can choose to enter a database user name and password, or simply clear the previously used database credentials.

The cockpit encrypts and stores the credentials, and allows you to connect to the database.

Related Information

Database Details
Search, Sort, and Filter Tools in the Resource Directory
Database Default Views [page 124]

4.2 Configurations and Configuration Templates

Compare the configuration parameters for managed databases or use a configuration template to capture a set of parameter values and apply them to other databases, systems, or hosts.

Related Information

Taking a Snapshot of a Database's Configuration [page 99]
Comparing Database Configurations [page 99]
Applying a Configuration Template [page 101]
Creating a Configuration Template [page 100]
Modifying a Configuration Template [page 103]
Deleting a Configuration Template [page 103]

4.2.1 Taking a Snapshot of a Database's Configuration

Save a configuration snapshot: a timestamped copy of a managed database's full set of configuration parameters.

Prerequisites

- Register the database whose configuration you want to capture.

Context

Snapshots let you capture an accurate record of each database's configuration, track configuration changes, and provide context to the historical data SAP HANA cockpit collects.

Procedure

1. On your database's *Database Overview* page, click the *Manage system configuration* link on the *Database Administration* card.
2. Click *Take Snapshot*.
3. (Optional) Enter a description for the snapshot. The cockpit automatically associates the snapshot with its database.
4. To see previous snapshots, click the *Snapshots* tab.

   The cockpit lists any previous snapshots of this database's configuration.
5. (Optional) To delete a snapshot, select it in the snapshots list and click ⊗*Delete Snapshot*.

4.2.2 Comparing Database Configurations

Compare the current configurations of two databases, compare two snapshots, or compare a current database configuration to a snapshot.

Prerequisites

- Add the databases whose configuration you want to compare.
- Ensure that the database or databases whose configurations you want to compare support configuration management. SAP HANA databases support configuration management in version SAP HANA 1.0 SPS 12 and later.

Context

To run a comparison, you select a source and a target to compare.

- When you compare two current configurations, the source and target databases must be of the same type (two SAP HANA databases, for example) and must be running the same software version.
- When you compare a current configuration to a snapshot, they must belong to the same database. That is, you can't compare a current configuration to a snapshot of another database.
- When you compare two snapshots, they must belong to the same databases. That is, you can't compare a snapshot of one database to a snapshot of another database.

Procedure

1. On your database's *Database Overview* page, click the *Manage database configuration* link on the *Database Administration* card.
2. Click *Compare*.
3. Choose a *Source*, choose the type of a target to compare it to (a *Snapshot* or a *Tenant Database*), and choose the *Target*.

   Snapshots are listed by timestamp. Hover your mouse over the timestamp in the list to see that snapshot's description (if it has one).
4. (Optional) To display in the results only parameters whose values are different in the source and the target, select the *Show differences only* checkbox.

   The cockpit filters out of the results those parameters whose values are the same in the source and the target. If there are no differences between the source and the target in the current configuration file, the cockpit displays `No differences found` in the parameter table.
5. To compare two snapshots:
   a. With the database selected in the left pane, click the *Snapshots* tab in the middle of the screen.
   b. Select the source snapshot from the list.
   c. Click *Compare*.
   d. On the *Database Configuration* screen, select the target snapshot from the *Target* dropdown list.

      Snapshots are shown by timestamp. Hover your mouse over the timestamp in the list to see that snapshot's description (if it has one).

4.2.3 Creating a Configuration Template

Set up a configuration template you can use to set parameter values on selected databases.

Prerequisites

You have the `Cockpit Template Administrator role`.

Context

Every configuration template has a layer: database, host, or system. When you apply a database-layer template, it affects only the database you select. When you apply a host-layer template, it affects all the databases on that host, and when you apply a system-layer template, it affects all the associated tenant databases.

Templates are created by setting the configuration parameters on one of your databases, then using that database to create the template.

Procedure

1. On the SAP HANA Cockpit *Home* page, click *Configuration Templates*.
2. Click *Create Template*.
3. Enter a name for your new template and a description. The description is optional.
4. Select the layer at which you want to create the template: database, host, or system.
5. Select the database, host, or system whose parameters serve as the model for the template. (You can add, remove, and reset the parameters later.)
6. Do one of:
   - Click the box at the top of the list to select all the parameters.
   - Click boxes to select one or more parameters individually.
7. Click *Review*.
8. Check the template and use the *Edit* links to correct it as needed.
9. Click *Create Template*.
10. (Optional) You can create a template from the *System Configuration* page of your cockpit database, but your template can only be created based on the host, system, or database that you're currently connected to. The steps are similar to the previous steps.

4.2.4 Applying a Configuration Template

Use a configuration template to set parameter values for databases, systems, or hosts.

Prerequisites

- You have the system privilege INIFILE ADMIN on the database to which you're applying a configuration template.
- You have permission to set configuration parameters on the target systems.

Context

Every configuration template has a layer: database, system, or host. When you apply a database-layer template, it affects only the database you select. When you apply a system-layer template, it affects all the associated tenant databases, and when you apply a host-layer template, it affects all the databases on that host.

Procedure

1. On your database's *Database Overview* page, click the *Manage database configuration* link on the *Database Administration* card.
2. Click *Configuration Templates*.
3. Select a configuration template from the list.

   Select the template's row in the list (not its radio button) to display its parameters and their values.
4. Click *Apply to Databases*.
5. (Optional) In the *Select Databases* window, enter a full or partial database or host name in the Search field to filter the list of databases.
6. Do one of:
   - Click the *Database*. A user with the Cockpit Template Administrator role has created at least one box at the top of the list to select all the database.
   - Click the boxes to select one or more databases individually.

   > **i Note**
   >
   > The template application fails if you select a database for which you lack either the privileges needed to change the configuration or valid credentials. You can enter credentials in the *Database Directory*.
7. Click *Apply Template*.

Results

A pop-up tells you whether the template was applied successfully and to how many databases it was applied.

Related Information

Creating a Configuration Template [page 100]

4.2.5 Modifying a Configuration Template

Rename a configuration template, change its description, add or remove parameters, or change parameter values.

Prerequisites

You have the `Cockpit Template Administrator role`.

Procedure

1. On the SAP HANA Cockpit *Home* page, click *Configuration Templates*.
2. Select the row (not the radio button) of the template you want to modify.
3. To change the template's name or description, click *Rename* and enter the changes.
4. To change a parameter's value, click its *Edit Specific Value* link, then:
   - Enter a new value and click *Save*, or
   - Click *Assign Default* and click *Save*, or
   - Click *Cancel*.
5. To remove a parameter from the template, click its ⊗ *(Delete Parameter)* icon.
6. To add a parameter, click *Include More Parameters* and select from the list.

   The parameters shown are parameters of the same layer as the one specified in the template and are from the database that was originally used to create the template.

4.2.6 Deleting a Configuration Template

Delete a configuration template from SAP HANA cockpit.

Prerequisites

You have the `Cockpit Template Administrator role`.

Context

Procedure

1. On the SAP HANA Cockpit *Home* page, click *Configuration Templates*.
2. Select the row (not the radio button) of the template you want to delete.
3. Click *Delete* and confirm that you want to delete the template.

4.3 SAP EarlyWatch Alert Service

Cockpit users with Support Hub user credentials and the required authorizations can configure *SAP EarlyWatch Alert* (EWA) in the Cockpit Manager to use the Solution Finder for SAP EarlyWatch Alert and the Alerts card inside SAP HANA cockpit.

The EWA functionality inside the SAP HANA cockpit collects diagnostic information from your cockpit system and sends it to SAP; SAP then sends back alerts and recommendations for the connected HANA databases in your cockpit system.

The Solution Finder for SAP EarlyWatch Alert offers a full text search of all SAP EarlyWatch Alert reports for any affected system. The search results display related alerts (sorted by severity), recommendations by SAP, and additional information. Alerts and recommendations sent back by SAP can also be viewed by all SAP HANA cockpit users from the *Alerts* card of your cockpit database.

4.3.1 Specifying Authorization for the Technical User

In order for the SAP HANA cockpit to collect and send information to the SAP EarlyWatch Alert service, each database registered through the cockpit requires a technical user with specific privileges.

Prerequisites

You must have the USER_ADMIN system privilege.

Context

SAP recommends that you set up a dedicated user account for each technical user. This user account shouldn't be used by people, but rather allocated for the purpose of connecting the database and the cockpit.

Set up SAP EarlyWatch Alert service for a specific HANA database in your SAP HANA cockpit. Set up the service by granting the following authorizations to the technical user of that database.

Procedure

1. For each of the SAP HANA database you want to register through the cockpit, create a technical user account or modify the existing user account that the cockpit uses to collect monitoring data.

   > **i** Note
   >
   > Use SQL to create the technical user required to register a database through the SAP HANA cockpit and grant the minimum necessary authorizations:
   >
   > ```
   > CREATE USER <username> PASSWORD <password> NO FORCE_FIRST_PASSWORD_CHANGE
   > VALID UNTIL FOREVER;
   > GRANT CATALOG READ to <username>;
   > GRANT SELECT on SCHEMA _SYS_STATISTICS to <username>
   > ```

2. For the technical user on each registered database for which the cockpit collects data for the SAP EarlyWatch Alert service, assign the SELECT object privilege on the following views:
   - change_entries
   - package_catalog
   - changes
   - active_object
   - inactive_object
3. Ensure that the technical user has the following:
   - Role: PUBLIC to access the Monitoring views.
   - Role: MONITORING to access the Statistics server views.

4.3.2 Managing SAP EarlyWatch Alert Global and Database-Specific Settings

In order for the SAP HANA cockpit to collect and receive data for the SAP EarlyWatch Alert service, you need to enter credentials for a valid global or specific S-User and configure transmission and collection settings.

Context

You can set a global S-User, whose credentials are used for all the databases in your cockpit. However, if you have multiple S-Users and want to use a specific S-User for a specific cockpit database, you can override the global S-User settings for that database. Overriding the global S-User can be helpful, for example, if you have a specific S-User attached to a specific customer and want to use those credentials for that customer's cockpit database.

> **i** Note
>
> SAP recommends setting the SAP EarlyWatch Alert service per cockpit database, rather than globally, specifically if your cockpit contains both development and production systems. The SAP EarlyWatch Alert service intended only for production systems.

Procedure

1. Connect to the Cockpit Manager and sign in as a cockpit administrator.
   Access the Cockpit Manager by following the *Manage Cockpit* link in the cockpit.
2. Choose one of the following options:

Option	Action
Manage global EWA settings	1. Select *Cockpit Settings* > *SAP EarlyWatch Alert* > . 2. Click *Edit*. 3. Enter the S-User credentials that you want to use for all databases in the cockpit. For instructions on how to create an S-User, see *SAP Note 2174416: Creation and activation of users in the Technical Users application - SAP ONE Support Launchpad*. 4. Select the check box to receive SAP EarlyWatch Alerts and send any collected data to the SAP EarlyWatch Alert service. **i Note** If this SAP HANA database is connected to an SAP Solution Manager or any other ABAP stack, you can receive an SAP EarlyWatch Alert even if you deselect the check box. However, if you deselect the check box labeled *Gather data to send to SAP EarlyWatch Alert service for analysis* and the database doesn't have a connection with an SAP Solution Manager or other ABAP stack, you don't receive an SAP EarlyWatch Alert. 5. Specify time of day at which data is transmitted daily from the cockpit to the SAP EarlyWatch Alert service. 6. Specify the day of the week and time at which data is collected from the registered database in the cockpit, and sent to the SAP EarlyWatch Alert service. 7. (Optional) Change the URL location where data is stored for transmission by both the SAP HANA cockpit and the SAP EarlyWatch Alert service. SAP recommends that you keep the default URL unless otherwise advised by an SAP representative. 8. (Optional) Add information for one or more SAP router instances. If more than one SAProuter instance is specified, multiple hops are used to transmit the data from the cockpit to SAP. 9. Click *Save*.

Option	Action
Manage database-specific settings	1. In the *Database Directory*, click *Manage Databases*. 2. Select the database that you want to set a database-specific S-User for and click *Edit*. 3. Enter your technical user credentials and click *Grant SAP EarlyWatch Alert permission*. 4. Scroll down to the *EarlyWatch Alert Support Hub User* section. 5. Enter the credentials of the S-User that you want to use for this specific database. > **i Note** > Setting a database-specific S-User overrides the global S-User settings. 6. (Optional) Override the global EarlyWatch Alert transmission and collection settings.

Related Information

SAP Note 2174416

4.3.3 Viewing SAP EarlyWatch Alerts

Cockpit users with Support Hub user credentials and the required authorizations can select the *SAP EarlyWatch Alert* tile to launch the Solution Finder for SAP EarlyWatch Alert in the SAP ONE Support Launchpad. Cockpit users without Support Hub user credentials can also view alerts and recommendations generated by SAP EarlyWatch Alert if it has been configured for their cockpit database using the *Alerts* card in the SAP HANA cockpit.

Prerequisites

SAP EarlyWatch Alert (EWA) has been configured for your cockpit database.

Context

The Solution Finder for SAP EarlyWatch Alert offers a full text search of all SAP EarlyWatch Alert reports for any affected system. The search results display related alerts (sorted by severity), recommendations by SAP and

additional information. You can select specific systems to see the complete corresponding paragraph from the relevant SAP EarlyWatch Alert report.

Procedure

1. Log on to the SAP HANA cockpit.
2. If you have Support Hub user credentials, you can view alerts, critical info, and recommendations generated by the SAP EarlyWatch Alert for all configured cockpit databases. View the information by selecting the *SAP EarlyWatch Alert* tile on the cockpit home page.
3. If you don't have Support Hub user credentials, you can view alerts and critical info generated by the SAP EarlyWatch Alert service for your cockpit database by clicking the *Alerts* card on your database's *Database Overview*.

   If a displayed alert is an EWA alert, the *Source* entry is *SAP EarlyWatch Alert service*. Only a selected subset of EWA alerts shown in the Solution Finder are forwarded to the SAP HANA cockpit.

5 Using the Database Overview Page to Manage a Database

The *Database Overview* page displays important metrics and available functions, regardless of whether the database you're managing is a single-host or multi-host system, or a tenant database.

Through the *Database Overview* page, you can view key health indicators for this specific database, such as database status, alerts, and database utilization. You also have access to tools that allow you to perform database administrations tasks, such as performance analysis, and executing SQL statements. Different parts of a single card can link to different views or applications. This way, you can see various components in a single view and decide whether to further examine issues by drilling down.

The *Database Overview* page is organized by cards, each of which represents a task or group of actions. You can search for a specific card, or filter by area to see related cards.

To launch the *Database Overview*, drill down on the name of the database from the *Database Directory*. Unless your administrator has enabled single sign-on, you'll need to connect to the database with a database user that has the system privilege CATALOG READ and SELECT on _SYS_STATISTICS.

> **i** Note
>
> If you register an offline database and display its *Database Overview*, much of the information that appears on the *Database Overview* of an online database is missing. You can use the cockpit to start the database: click the *Stopped* icon, then click *Start System*.

The cards that appear on the *Database Overview* depend on whether you're connected to a tenant or system database and whether or not it's part of a multi-host system.

Services

Your status can be running, running with issues, or stopped. Clicking on this status brings you to *Manage Services* where you can stop or kill a service, and start or stop a system.

Alerts

Alert counts for the database are displayed for high- and medium-priority alerts, broken down by the nine alert categories defined in SAP HANA. (You can refresh the displayed data by using the manual or auto-refresh icons in the top-right corner). Clicking on the *Alerts* card brings you to the *Alert Monitor* for the database. In the bottom-right corner, there's a status message showing vital information about SAP HANA processes that collect data. By noting the status messages within the card, you can easily find out the validity of what you are seeing.

Usage and Performance Metrics

You can monitor key database metrics through the *CPU Usage*, *Memory Usage* and *Disk Usage* cards, as well as the *Threads*, *Sessions*, and *Monitor Statements* cards. In a multi-host system, each host is represented by a clickable bar, with the selected host having a time graph displayed to the right of the bar chart. Hover over the bars to see details for the selected host. If a bar is highlighted, there's an associated high (red) or medium (yellow) alert. With single-host databases, since there's only one host, no bar graphs are displayed. By viewing this high-level information, you can decide whether to drill down to the *Performance Monitor*. See *Monitoring and Analyzing with the Performance Monitor*.

Smart Data Access

Without copying data directly into an SAP HANA database, you can use Smart Data Access to access remote data as if it were stored in local tables. Refer to the Smart Data Access section of the *SAP HANA Administration Guide*.

System Replication

If a database is part of a system replication configuration, you can monitor the status of replication between the primary system and one or more secondary systems. See *Monitoring SAP HANA System Replication with the SAP HANA Cockpit*.

Additional Functionality

You can launch additional functionality by selecting any of the links organized under the headings *Monitoring*, *DB Administration*, *User and Role Management*, *Alerting and Diagnostics*, *Other Administration*, *Application Lifecycle Management*, *Platform Lifecycle Management*, and *Help*. Specific links and related tasks are described in the subsequent topics of this guide.

Security

The *Data Storage Security*, *Auditing*, *Authentication* blocks, and the *Security Related Links* help you to monitor many critical security settings. Additionally, you can perform administration tasks related to data and communication encryption, and audit logging. See *Monitoring Critical Security Settings*.

Performance Management

Use *Analyze Workload*, *Capture Workload*, and *Replay Workload* to manage performance. See *Capturing and Replaying Workloads*.

Use *Analyze Workload* to analyze the database performance. See *Monitoring, Analyzing, and Improving Performance*.

Additionally Installed SAP HANA Contexts

Other cockpit features that allow you to manage additionally installed contexts (for example, SAP HANA dynamic tiering) are only visible and available if the specific context has been installed.

Related Information

Monitoring Tenant Databases in SAP HANA Cockpit [page 418]
Services [page 159]
Alerts [page 126]
The Performance Monitor [page 515]
Monitoring, Analyzing, and Improving Performance [page 513]
Monitoring SAP HANA System Replication [page 167]
Monitoring Critical Security Settings [page 221]
Capturing and Replaying Workloads [page 532]
Connect to a Database With SSO or SAP HANA Credentials [page 97]
Personalizing the Database Overview Page [page 21]

5.1 Authorizations Needed for Monitoring and Administration

To view information about the SAP HANA database and access the various applications for administration and monitoring, you need to connect to the database with a database user with appropriate database privileges. If you don't have the required privileges, specific tiles, features, or actions aren't available to you.

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

Database Monitoring and Administration

These tables list the database privileges required to view information about an SAP HANA database on the *Database Overview* page and to access monitoring and administration functions on subsequent pages.

Monitoring

To Access...	You Need These SAP HANA Privileges...

Services	• CATALOG READ system privilege
	• SELECT privileges on:
	  ○ SYS_DATABASES.m_services
	  ○ SYS_DATABASES.m_service_memory
	  ○ SYS_DATABASES.m_service_statistics
	  ○ SYS_DATABASES.m_heap_memory_reset
	  ○ SYS.m_services
	  ○ SYS.m_service_memory
	  ○ SYS.m_service_statistics
	  ○ SYS.m_heap_memory_reset
	  ○ _SYS_STATISTICS.STATISTICS_SCHEDULE
	  ○ _SYS_STATISTICS.HELPER_ALERT_CHECK_INACTIVE_SERVICES_AGE

Alerts	SELECT privileges on _SYS_STATISTICS
Memory Usage	No additional authorization required
Memory Analysis	• CATALOG READ system privilege • SELECT on _SYS_STATISTICS object privileges
CPU Usage	No additional authorization required
Disk Usage	No additional authorization required
Performance Monitor	No additional authorization required
Monitor Statements	To enable or disable memory tracking, you need the INIFILE ADMIN privilege. To cancel the session, you need the SESSIONS ADMIN privilege.
Sessions	To cancel sessions or operations, you need the SESSION ADMIN privilege.
Threads	To cancel sessions or operations, you need the SESSION ADMIN privilege.
Monitor expensive statements	To configure expensive statements, you need the INIFILE ADMIN privilege.
Open SQL plan cache	To configure SQL plan cache, you need the INIFILE ADMIN privilege.
Open blocked transactions	No additional authorization required
Smart Data Access	No additional authorization required

Alerting and Diagnostics

To Access...	You Need These SAP HANA Privileges...
Configure alerts	INSERT, EXECUTE, DELETE, and UPDATE privileges on _SYS_STATISTICS
Plan trace	• To turn plan trace on or off, you need the TRACE ADMIN privilege. • To retrieve the results, you need the SELECT privilege on SYS schema. The SYSTEM user or a user with the SAP_INTERNAL_HANA_SUPPORT_ROLE has this authorization.

System Replication

To Access...	You Need These SAP HANA Privileges...
System Replication	No additional authorization required

Other Administration

To Access...	You Need These SAP HANA Privileges...
Apply configuration templates	System privilege INIFILE ADMIN

Application and Platform Lifecycle Management

To Access...	You Need These SAP HANA Privileges...
Platform Lifecycle Management	`<sid>adm` privileges
Application Lifecycle Management	Not applicable. For more information about the availability of the application lifecycle management GUI, see the section on SAP HANA application lifecycle management in the *SAP HANA Administration Guide*.

System Information

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

Security

Security

To Access...	You Need These SAP HANA Privileges...
Data Encryption	• To enable/disable encryption: ENCRYPTION ROOT KEY ADMIN • To view SSFS master key information: RESOURCE ADMIN
Auditing	To see, create, change, and delete audit policies, system privilege AUDIT ADMIN. (AUDIT ADMIN doesn't allow you to delete the audit log.) To delete the audit log, AUDIT OPERATOR. To change the auditing status, one of the following system privileges: • AUDIT ADMIN • INIFILE ADMIN To see the system views AUDIT_LOG, system privilege AUDIT READ, AUDIT ADMIN, or AUDIT OPERATOR. To see the audit policies, the system privilege CATALOG READ or AUDIT ADMIN.
Authentication	To see information about authentication, system privilege INIFILE ADMIN. To be able to see the password blacklist on opening the *Password Policy and Blacklist* page, you need SELECT privilege on _SYS_PASSWORD_BLACKLIST (_SYS_SECURITY).

User & Role Management

To Access...	You Need These SAP HANA Privileges...
Manage users	To view, create, and manage users, system privilege USER ADMIN. To view users, system privilege CATALOG READ. To change your own password, no privilege is required.
Assign roles to users	To view and assign roles, system privilege ROLE ADMIN. To view roles, system privilege CATALOG READ.
Assign privileges to users	To assign privileges, the privileges must be granted with the permission to grant them to other users. To view all the privileges granted, system privilege CATALOG READ.

To Access...	You Need These SAP HANA Privileges...
Manage roles	To create and manage roles, system privilege ROLE ADMIN and the privileges required to grant specific privileges to roles. To view roles, system privilege CATALOG READ. To grant object privileges on your own objects to a role, no privilege is required.
Manage user groups	To create and modify user groups, system privilege USER ADMIN. To modify an existing user group configured for exclusive administration, the object privilege USERGROUP OPERATOR on the user group.

Security-Related Links

To Access...	You Need These SAP HANA Privileges...
Manage certificates	To read and change certificates, CERTIFICATE ADMIN. To read certificates, CATALOG READ or TRUST ADMIN.

Manage certificate collections	To create certificate collections and assign certificates, TRUST ADMIN.
	
	To see, change, or delete a certificate collection, you must:
	
	• Have the system privilege DATABASE ADMIN and the object privilege REFERENCES, DROP, or ALTER on a certificate collection.
	• Be the owner of the certificate collection and have one of the following system privileges:
	  ○ SSL ADMIN
	  ○ USER ADMIN
	  ○ CERTIFICATE ADMIN
	
	To see certificate collections, CATALOG READ.

Network security information	System privilege CATALOG READ
View anonymization report	System privilege CATALOG READ
Manage SAML identity providers	System privilege USER ADMIN
Manage JWT identity providers	System privilege USER ADMIN
Security administration help	No additional authorization required
SAP HANA security website	No additional authorization required
Security checklists	No additional authorization required

Backup and Recovery (SAP HANA Cockpit)

Authorizations for Backup and Recovery (SAP HANA Cockpit)

To Perform This Task...	You Need These SAP HANA Privileges...		
	System Database	Tenant Database	Tenant Database (Through the System Database)

View, create, and cancel database backups	`BACKUP ADMIN` or `BACKUP OPERATOR` (recommended for batch users only).	`BACKUP ADMIN` or `BACKUP OPERATOR` (recommended for batch users only).	`DATABASE BACKUP OPERATOR`
			`DATABASE BACKUP ADMIN`
			`DATABASE ADMIN`

Delete database backups and backups of the backup catalog	`BACKUP ADMIN`	`BACKUP ADMIN`	`DATABASE BACKUP ADMIN` `DATABASE ADMIN`
Recover or copy a database	Operating system user `<sid>adm`	(not possible)	`DATABASE RECOVERY OPERATOR` `DATABASE ADMIN`

Schedule backups	`BACKUP ADMIN`
	`SELECT`, `UPDATE`, and `DELETE` privileges for the following tables:
	• `_SYS_XS.JOB_SCHEDULES`
	• `_SYS_XS.JOBS`
		`BACKUP ADMIN`
		`SELECT`, `UPDATE`, and `DELETE` privileges for the following tables:
		• `_SYS_XS.JOB_SCHEDULES`
		• `_SYS_XS.JOBS`
			`DATABASE BACKUP OPERATOR`
			`DATABASE BACKUP ADMIN`
			`DATABASE ADMIN` and SELECT, UPDATE, and DELETE privileges for the following tables:
			• `_SYS_XS.JOB_SCHEDULES`
			• `_SYS_XS.JOBS`

To Perform This Task...	You Need These SAP HANA Privileges...		
	System Database	Tenant Database	Tenant Database (Through the System Database)
Configure backups	To display the backup configuration settings, `BACKUP ADMIN`. To configure **backup retention**, you also need: • `SELECT`, `UPDATE`, and `DELETE` privileges for table `_SYS_XS.JOB_SCHEDULES` • `SELECT privileges for table _SYS_XS.JOBS`	`DATABASE BACKUP ADMIN` To display the backup configuration settings, `BACKUP ADMIN`. To configure **backup retention**, you also need: • `SELECT`, `UPDATE`, and `DELETE` privileges for table `_SYS_XS.JOB_SCHEDULES` • SELECT privileges for table `_SYS_XS.JOBS`	`DATABASE BACKUP ADMIN`

System-wide Backup Configuration

To change default configuration settings for all the tenant databases, you need `DATABASE ADMIN`.

For more information, see *Authorizations Needed for Backup and Recovery* in the *SAP HANA Database Backup and Recovery* section of the *SAP HANA Administration Guide*.

Performance Management

Workload Management

To Access...	You Need These SAP HANA Privileges...
Analyze Workload (Based on Thread Samples)	System privileges CATALOG READ and INIFILE ADMIN.
Analyze Workload (Based on Engine Instrumentation)	System privilege WORKLOAD ANALYZE ADMIN.

To Access...	You Need These SAP HANA Privileges...
Capture Workload	Workload Capture Admin used to start captures, trigger backups, and collect Abstract SQL Plans in the source system. Privileges: • Requires the WORKLOAD CAPTURE ADMIN privilege to capture workloads • It can have the BACKUP OPERATOR privilege to trigger backups • It can have the INIFILE ADMIN privilege to see the previously used optional filters on the capture configuration page
Replay Workload	The Control Replay Admin used for preprocessing and replaying in the control system. Privileges: • Requires the WORKLOAD REPLAY ADMIN privilege for preprocessing and replaying workloads, as well as generating replay reports • Requires the WORKLOAD ANALYZE ADMIN privilege for loading and analyzing workloads • Requires the CATALOG READ privilege for generating replay reports The Target Replay Admin used to execute replays and reset user passwords in the target system. Privileges: • Requires the WORKLOAD REPLAY ADMIN privilege to execute replays

Related Information

Authorizations Needed for Backup and Recovery

5.2 Cards Available on the Database Overview Page

The SAP HANA cockpit offers many pages to help you administer and monitor databases. These pages are accessible by selecting individual cards or links on the *Database Overview* page. If you don't have the required privileges, specific cards, features, or actions aren't available to you.

Monitoring View

Card	Description
Services	Monitors the status services. Opens the *Manage Services* page where you can view database usage of individual database services, as well as perform other administration tasks such as stopping services. Indicates overall database health and the number of services running and not running are indicated. If the database is distributed across multiple hosts, it includes all services on worker hosts.
Alerts	Monitors the status of the database and its services and the consumption of system databases. Indicates the number of high and medium alerts currently raised in the database. Opens the *Alerts* page where you can view and analyze alert details.
Memory Usage	Monitors the total amount of memory currently used by the SAP HANA database in relation to the allocation limit. *Performance Monitor* page - Lets you visualize and explore the usage history of key system databases (CPU, memory, and disk). When accessed from the *Memory Usage* card, memory-related KPIs are automatically selected. *Analyze Workloads* page - Monitors overall system health and lets you identify the root cause of potential performance issues. *(More)* - Additional pages available for monitoring cache, memory, and load configuration.
CPU Usage	Monitors the percentage of CPU used. 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. *Performance Monitor* - Lets you visualize and explore the usage history of key system databases (CPU, memory, and disk). When accessed from the *CPU Usage* card, CPU-related KPIs are automatically selected. *Analyze Workloads* - Monitors overall database health and lets you identify the root cause of potential performance issues.

Card	Description
Disk Usage	Monitors the **total** usage of all disks, including space used by non-SAP HANA data. The disk with the highest (most critical) disk usage is shown in more detail. *Performance Monitor* - Lets you visualize and explore the usage history of key system databases (CPU, memory, and disk). When accessed from the *Disk Usage* card, disk-related KPIs are automatically selected. *Monitor Disk Volume* - Monitors disk statistics to ensure that there's enough space on disk for data volumes and log volumes. *Analyze Workloads* - Monitors overall database health and lets you identify the root cause of potential performance issues.
SQL Statements	Monitors and analyzes different types of statements in your database, including active, SQL Plan cache, and expensive statements. Lets you enable statement memory tracking.
Monitoring	*Performance Monitor* - Lets you visualize and explore the usage history of key system databases (CPU, memory, and disk). *Monitor Table Usage* - Monitors the number of hosts monitored and the name of the host with the highest memory usage. Opens *Table Usage* page, where you can visualize tables by size, explore the usage history of tables, and move tables to warm storage. *Open blocked transactions* - Monitor transactionally blocked threads. *Profile memory* - Analyzes the memory consumption of the SAP HANA database over time can help you pinpoint bottlenecks, identify patterns, and forecast requirements.
Threads	Monitors the number of currently active and blocked threads.
Admission Control	Manages peak load by applying processing limits and determining how to handle new requests if the database is close to the point of saturation.
Sessions	Analyzes the sessions connected to your SAP HANA database helps you identify which applications or which users are currently connected to your database, as well as what they are doing in terms of SQL execution.
System Replication	Monitors status of system replication. On SYSTEMDB, the card lets you configure the system for replication

Security and User Management View

Card	Description
Data Encryption	Configures SAP HANA encryption services, including data volume and redo logs and data and log backups.
Auditing	Monitors and configure audit policies, audit trails, and auditing-related alerts.

Card	Description
Anonymization Report	Monitors analytics on data while still protecting the privacy of individuals through anonymization views in the case of SQL, and as anonymization nodes in the case of calculation views.
User and Role Management	• *User Management* - Monitors and manages database users. • *Role Assignment* - Monitors and manages roles granted to database users. • *Privilege Assignment* - Monitors and manages privileges granted to database users. • *Role Management* - Monitors and manages roles to be granted to database users. • *User Group Management* - Monitors and manages user groups for database users. • *Authorization Dependency Viewer* - Analyzes database objects and their dependencies to help resolve object authorization errors..
Security Related Links	• *Certificate Store* - Monitors the in-database repository for X.509 client certificates. • *Certificate Collections* - Monitors and manages certificate collections. • *View network security information* - Monitors the cryptographic library in use in the database and the minimum accepted version of the transport layer security/secure sockets layer (TLS/SSL) protocol and lets you can see more detailed information about network configuration. • *SAML Identity Providers* - Monitors and manages Security Assertion Markup Language (SAML) identity providers. • *JWT Identity Providers* - Monitors and manages JSON Web Tokens (JWT) identity providers. • *Security Checklist* • *Security Administration Help* - View documentation describing those security administration tasks that you can perform using the SAP HANA cockpit. • *SAP HANA Security Website* - Link to the SAP HANA security website.
Authentication	Manages password policies (default or customized) and password blacklists.
Insufficient Privilege Details	Displays information about an "insufficient privilege" error.

Administration View

Card	Description
Capture Workload	Monitors captured workloads. *Start new Capture* - Opens the *Capture Configuration* page, where you can configure new workloads for capture.
Replay Workload	Monitors the number of replayed workloads. Opens the *Replay Workload* page, where you can preprocess workloads, replay preprocessed workloads and monitor during workload replay. *Start New Replay* - Starts and stops any started replays.

Card	Description
Database Backups	Manages backup and recovery of databases. Opens the *Backup Catalog* page, where you can view a list of existing backups. *Backup Schedules* - Monitors and manages scheduled backups. *Recover Databases* - Recovers a database by selecting from a list of existing backups. *Copy Database* - Copies an SAP HANA database by recovering it to the same database or to a different database. *Backup Configuration* - Monitors and manages backup configurations.
Smart Data Integration	See Smart Data Integration Guide (https://help.sap.com/viewer/p/HANA_SMART_DATA_INTEGRATION).
Database Administration	*Manage database configuration* - Monitors and manages configuration (*.ini) files at the system, database, and host level. *Database Explorer* - Access the to query information about the database, view information about your database's catalog objects, and open the SQL console to execute SQL statements. *Manage workload classes* - Monitor and manage workload classes. *Manage database licenses* - Monitor and manage license keys for the database.
Manage SQL Performance	Analyzing statement performance to understand performance issues of a query execution and other query execution aspects of the SAP HANA database. Tools include Statement Hints, SQL Plan Stability, Plan Trace, Saved Plans, and Data Cache.
Recommendations	Monitor recommendations to get suggestions for changes you can apply to the SAP HANA database to increase its performance and operation. Recommendation type s include Ad Hoc, Physical Design, and SQL.
Alerting and Diagnostics	• *Alerts Definitions* - Opens *Alert Definitions* where you can configure alert schedules and thresholds and set up e-mail notification. • *Configure Tracing* • *View trace and diagnostic files* - Opens the *Trace* tool of the SAP HANA Database Explorer for SAP HANA cockpit. **i Note** The *Trace* tool opens in a new window. • Manage full system information dumps.
Database Information	View general information about the SAP HANA database, such as operational status and database version, can assist you to monitor your database.

Card	Description
Table Distribution (Tenant only)	Monitor and save the current table distribution, automatically generate an optimized table distribution, re-run a previously executed plan, or restore a saved plan. *View Current Table Distribution* - View how tables are distributed across the hosts.
Other Administration	*Manage Hadoop Cluster* - Navigate to the Apache Ambari Web site and monitor Hadoop clusters.
Smart Data Access	• *Running Statements* - Monitors the number of running remote statements. Opens the *Remote Statements Monitor* page, where you can analyze remote statements in the database. • *Active Connections* - Monitors the number of active remote connections. Opens the *Remote Connections Monitor* page, where you can analyze remote connections in the database.
Platform Lifecycle Management (SYSTEMDB only)	Specific links and related tasks that encompasses the installation and update of an SAP HANA server, mandatory components, and additional components, as well as the post-installation configuration.

6 Database Default Views

Track database health, services, memory allocation, performance, and alerts.

The *Database Overview* page contains a series of cards. Each card provides a starting point to a group of related applications that allow you to monitor and manage the database. Click the card title to open the underlying application for the card.

The cards are grouped into five views:

- Monitoring
- Security & User Management
- Administration
- All

Related Information

Using the Database Overview Page to Manage a Database [page 109]
Start a Database [page 164]
Stop a Database [page 166]
Services [page 159]
Memory Analysis [page 139]
Alerts [page 126]
Alert Configuration [page 130]
Managing Workload Classes [page 457]
Monitor Table Usage [page 530]
Database Configuration in SAP HANA Cockpit [page 443]
Monitoring Multi-Host Health [page 191]
Manage Hadoop Clusters [page 478]
Database Information [page 477]

6.1 Getting to the Database Overview Page

There several ways to get to the *Database Overview* page, but they all start from the SAP HANA Cockpit *Home* page.

Using the Database Directory Tile

The *Database Directory* tile takes you to the *Database Directory* page which contains all the databases register in the SAP HANA Cockpit that you have access to. Click the database name to display its *Database Overview* page.

Using a Group Tile

When databases are registered in SAP HANA Cockpit Manager, you can organize them into groups. So if you know the group of the database you want to use, click the group's tile. Only those databases assigned to the group that you have access to appear on the *Database Directory*. Click the database name to display its *Database Overview* page.

Using a Shortcut Tile

Once you get to the *Database Overview* page of a database, you can create a shortcut to the page that appears on the *Home* page. Clicking the shortcut bypasses the *Database Directory* page, taking you directly to the *Database Overview*.

7 Monitoring View

This view helps you to monitor many details pertaining to single- and multi-host resources.

You can perform such tasks as:

- Monitor overall database health
- Monitor status and resource usage of individual database services
- Analyze database performance across a range of key performance indicators related to memory, disk, and CPU usage
- Analyze the comparative memory utilization of column tables
- Analyze the memory statistics of the components of database services
- Monitor alerts occurring in the database and analyze patterns of occurrence
- Configure the alerting mechanism, for example, change alert threshold values, switch alert checkers on/off, and check for alerts out of schedule
- Monitor the status of system replication (if enabled)

It's important that you monitor the operation of SAP HANA databases on a regular basis. Although SAP HANA actively alerts you of critical situations, keeping an eye on resource usage and performance will help you identify patterns, forecast requirements, and recognize when something is wrong.

7.1 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 to 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, it 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 the configured threshold values. For example, if 90% of available disk space has been used, a low priority alert is issued; if 98% has been 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* in the *SAP HANA Administration Guide*.

For information on configuring alerts in SAP HANA cockpit, see *Configuring Alerts*.

Alerts are organized by category or KPAs (key performance areas), which is a predefined collection of the alert categories.

KPA	Alert Categories
Availability	• Availability • Backup • Diagnosis Files • Security
Performance	• CPU • Disk • Memory
Capacity	• Configuration • Sessions/Transactions

A summary of high and medium alerts in the database appears on the *Alert* card on the *Database Overview* page for a database. On the *Alters* card, you can view alerts by category or KPA. The category or KPA with the most high alerts is listed first, followed by medium alerts. When two categories/KPAs have the same number of high and medium alerts, the categories/KPAs are listed alphabetically. Categories/KPAs with no medium or high alerts aren't listed. For alert details on a specific category or KPA, click the row. For alert details on all alerts, click the title of the *Alerts* card.

If you're using the KPA view and you display alerts for a specific KPA, the resulting list on the *Manage Alerts* page includes alerts for all categories within the specified KPA. You can't filter, sort, or group the list on the*Manage Alerts* by KPA.

To view alerts, you must have the CATALOG_READ system privilege and the SELECT privilege on the _SYS_STATISTICS schema.

Related Information

Alert Configuration [page 130]
The Statistics Service
Search, Sort, and Filter Buttons [page 18]

7.1.1 Alert Summary and Details

On the *Manage Alerts* page, the default summary list includes all current alerts with a priority of high, medium, and error. You can filter, sort, and view details, as needed.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

- You have the CATALOG_READ system privilege and the SELECT privilege on the _SYS_STATISTICS schema.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, choose one of the following options:

Task	Action
Display alerts from a single category or KPA	Click the row of the category or KPA on the *Alerts* card.
Display all high and medium alerts	Click in the title area of the *Alerts* card.

A list of alerts appears on the *Manage Alerts* page.

Detail	Description
Priority	Indicates the severity of the alert and how quickly action needs to be taken. ○ Information - Action is recommended to improve system performance or stability. ○ Low - Medium-term action is required to mitigate the risk of downtime. ○ Medium - Short-term action is required (few hours, days) to mitigate the risk of downtime. ○ High - Immediate action is required to mitigate the risk of downtime, data loss, or data corruption. ○ Error - Immediate action is required to fix the issue. Use trace files to help track and resolve the issue.
Alert	Provides a definition of the alert.
Time	Indicates the time that the alert was triggered.
Alerting Host & Port	Provides the 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 identification allows you to ensure the availability of secondary systems by addressing issues before an actual failover. For more information about monitoring secondary systems in SAP HANA, see *Monitoring Secondary Sites* in the *SAP HANA Administration Guide*.
Source	Indicates where the alert originated from. The source could be the database itself, or a service such as the SAP EarlyWatch Alert service.

Detail	Description
Category	Indicates the category of the alert checker that issued the alert. Alert checkers are grouped into categories. The categories are: ○ Availability ○ Backup ○ Configuration ○ CPU ○ Diagnosis Files ○ Disk ○ Memory ○ Other ○ Security ○ Sessions/Transactions

2. Click an alert to see its definition and more details, such as past occurrences, proposed solutions, and next scheduled runs.

Related Information

Alert Configuration [page 130]
Monitoring Secondary Systems
Search, Sort, and Filter Buttons [page 18]

7.1.2 View Past Alerts

Current alerts are alerts that were triggered by the last scheduled run of the alert definition. Past alerts are the historical record of past current alerts.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the CATALOG_READ system privilege and the SELECT privilege on the _SYS_STATISTICS schema.

Context

The *Manage Alerts* page always displays the most recently triggered (current) alerts. When a scheduled alert definition runs, current alerts on the *Manage Alerts* page for the scheduled definition become past alerts,

disappearing from the list. New alerts for the definition triggered by the run appear a as current alerts on the list.

When Deactivating an alert definition, doesn't remove its most recently triggered alert from the list. To remove these alerts, reactivate the alert definition and either resolve the issue or adjust the threshold so the alert is no longer triggered.

If you deactivate the schedule for an alert that appears on the current list, the triggered alerts remain on the list, even if the triggering issue is resolved. The current alerts become past alerts when you reactivate the alert definition and allow it to run for at least one interval.

Filter the list using the type *Past*. To see the history of past alerts for a specific alert, display the alert details and select the frequency of the past occurrences.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click the *Alerts* card.
2. Choose one of the following:

Task	Action
For a list of past alerts for the database,	filter the list using the type *Past*. Define additional filters to further refine the alert list.
For a graphical display of past alerts for a single alert definition,	click the alert to display its details. Click *Occurrences* and set the past time frame.

7.1.3 Alert Configuration

Use the configuration options to tailor alerts in the SAP HANA database to your needs.

Alerts are based on predefined definitions, many of which are switched on by default. While you can't add new definitions, you can modify the predefined definitions by:

- Switching a specific alert definition on or off.
- Changing the threshold values that trigger alerts of different priorities, where applicable.
- Setting up e-mail notifications so that specific people are informed when alerts are issued.

7.1.3.1 Configure Alert 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've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the CATALOG_READ system privilege.
- You have SELECT, INSERT, DELETE, and UPDATE privileges on the _SYS_STATISTICS schema.

Context

Thresholds can be configured for any alert definition that measures variable values that stay within certain ranges. For example, the percentage of physical memory used or the age of the most recent data backup. Many alert definitions verify only whether a certain situation exists or not. Threshold values can't be configured for these alert definitions. For example, alert definitions 4 detects services restarts. If a service was restarted, an alert is issued.

You can set the threshold when creating a new alert or adjust the threshold of a triggered alert.

Procedure

1. To:

Task	Action
Adjust the threshold on a triggered alert.	1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click the *Alerts* card. 2. On the *Manage Alerts* page, click the triggered alert containing the thresholds you want to change. 3. Click *Edit Alert Definition*.
Configure the threshold on an alert definition.	1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Alert Definitions* on the *Alerts* card. 2. Click the alert definition containing the thresholds you want to change. 3. Click *Edit*.

2. Change the threshold values as required.
3. (Optional) To apply the alert definition to multiple databases, click *Apply to Other Databases*, select databases from the list, and click *OK*.

4. Save the alert definition.
5. (Optional) Click *Check Now* to run the alert definition immediately.

   A message appears letting you now if the modified alert definition triggered an alert.

Related Information

Assign Roles to Database Users [page 358]
Alert Summary and Details [page 127]

7.1.3.2 Switch Alerts Off/On

If you no longer want a particular alert to be issued, you can switch off the underlying alert definition. The system automatically switches off an alert definition when it fails to run, for example, due to a shortage of system resources.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the CATALOG_READ system privilege.
- You have SELECT, INSERT, DELETE, and UPDATE privileges on the _SYS_STATISTICS schema.

Context

In some situations, you may want to stop a particular alert from being issued, either because it's unnecessary (for example, alerts that notify you when there are other alerts in the system) or because it isn't 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 definition back on again at any time.

For a list of switched off alert definitions, see *Check Alert Definitions Status*.

Procedure

1. To:

Task	Action
Turn on/off an alert definition that has triggered an alert.	1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click the *Alerts* card. 2. On the *Manage Alerts* page, click the triggered alert to be switched on/off. 3. Click *Edit Alert Definition*.
Turn on/off an alert definition.	1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Alert Definitions* on the *Alerts* card. 2. Click the alert definition to be turned on/off. 3. Click *Edit*.

2. Set the *Schedule Active* switch to *No* or *Yes*.
3. (Optional) To apply the change to multiple databases, click *Apply to Other Databases*, select databases from the list, and click *OK*.
4. Save the alert definition.

Related Information

Check Alert Definition Status [page 135]

7.1.3.2.1 Disabling Alerts for a Table or Schema

It's possible to disable an alert for a particular table or schema.

Disabling alerts is supported for the alerts "Record count of non-partitioned column-store tables" (ID 17) and "Table growth of non-partitioned column-store tables" (ID 20).

To exclude an alert to be issued for a particular table, use the following SQL statement:

```
INSERT INTO _sys_statistics.statistics_exclude_tables VALUES (<alert_id>,
'<schema_name>', '<table_name>')
```

To exclude an alert to be issued for all tables of a particular schema, use the following SQL statement:

```
INSERT INTO _sys_statistics.statistics_exclude_tables VALUES (<alert_id>,
'<schema_name>', null)
```

To re-enable the alerts, delete the entries from the table `_sys_statistics.statistics_exclude_tables`.

7.1.3.3 Run Alert Definitions Manually

You can run alert definitions manually, outside the scheduled interval.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the CATALOG_READ system privilege.
- You have SELECT, INSERT, DELETE, and UPDATE privileges on the _SYS_STATISTICS schema.

Context

In some cases, you may want to check for a particular alert outside of its configured schedule. Running an alert definition on-demand doesn't change its next configured interval to run. An alert definition must have an active status to run it manually.

Procedure

1. To:

Task	Action
Manually run an alert definition that has a triggered an alert.	1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click the *Alerts* card. 2. On the *Manage Alerts* page, click the triggered alert you want to manually run. 3. Click *Edit Alert Definition*.
Manually run an alert definition.	1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Alert Definitions* on the *Alerts* card. 2. Click the alert definition you want to manually run. 3. Click *Edit*.

2. Click *Check Now*.

   A message appears letting you now if the definition triggered an alert.

Related Information

Open the SAP HANA Cockpit [page 13]

Alert Checker Details
Alert Checker Statuses
Switch Alerts Off/On [page 132]

7.1.3.4 Check Alert Definition Status

Check the current status of an alert definition.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the CATALOG_READ system privilege and the SELECT privilege on the _SYS_STATISTICS schema.

Context

By default, the *Define Alerts* page lists all alert definitions. Active definitions are green, and switched off definitions are yellow. Filter by status to focus the list.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Alert Definitions* on the *Alerts* card.
2. Click *View Alerts*.
3. (Optional) Set the Status filter.
4. (Optional) If the list is too long, filter by additional columns to further focus the list.

7.1.3.5 Set Up Email Notification Defaults

You can configure alert definitions so that you and other responsible administrators receive push notifications by email when alerts are issued.

You can configure one or more default recipients to be notified when any alert definition issues an alert. In addition, if different people need to be notified about different alerts, you can configure dedicated recipients for these alert definitions.

Note the following behavior:

- If you configure definition-specific recipients, default recipients are **not** notified.
- If you delete all definition-specific recipients, default recipients are notified again, if configured.
- You can configure definition-specific recipients regardless of whether or not default recipients are configured.

The configured recipients receive an email when an alert definition issues an alert. If the alert definition issues the same alert the next time it runs, no further email is sent. However, when the alert definition runs and it doesn't issue an alert, indicating that the issue is resolved or no longer occurring, a final email is sent

Related Information

Open the SAP HANA Cockpit [page 13]
Alert Checker Details
Alert Checker Statuses

7.1.3.5.1 Configure Default Sender Notification

Configure a sender to be used for email notification for alerts.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the CATALOG_READ system privilege and the SELECT privilege on the _SYS_STATISTICS schema.

Context

Once defined, the sender information is used for all alerts triggered for which email notification is defined.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Alert Definitions* on the *Alerts* card.
2. Click *Configure Email*, then *Sender*.

3. Enter:
    - The email address to be used as the sender.
    - The mail server that the system sends the emails to

      > **i** Note
      >
      > The statistics service doesn't support a mail server that requires additional authentication.

    - The default SMTP port is 25. If the configured mail server uses a different port, you must enter it.

7.1.3.5.2 Configure Default Recipient Email Notifications

Configure recipients to receive email notification for alerts.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the CATALOG_READ system privilege and the SELECT privilege on the _SYS_STATISTICS schema.

Context

Default recipients are notified about alerts generated by all alert definitions **except** those that have definition-specific recipients configured. The system validates the format of the email address but not the address itself.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Alert Definitions* on the *Alerts* card.
2. Click *Configure Email*, then *Default Recipient*.
3. Enter a recipient email address.
4. (Optional) Click *Add email* to define additional default recipients.
5. (Optional) Click *Send test email* to send a test email to the specified default recipient(s).
6. Click *Save*.

7.1.3.5.3 Configure Definition-Specific Email Notifications

Configure recipients to receive email notifications for alerts on a specific alert definition.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the CATALOG_READ system privilege and the SELECT privilege on the _SYS_STATISTICS schema.

Context

The default recipients for email notifications aren't notified of alerts for the specific alert definition when definition-specific recipients are defined.

Procedure

1. To:

Task	Action
Define email recipients for an alert definition that has triggered an alert.	1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click the *Alerts* card. 2. On the *Manage Alerts* page, click the triggered alert to be switched on/off. 3. Click *Edit Alert Definition*.
Define email recipients for an alert definition.	1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Alert Definitions* on the *Alerts* card. 2. Click the alert definition to be turned on/off. 3. Click *Edit*.

2. Enter a recipient email address.
3. (Optional) Click *Add email* to define additional default recipients.
4. Click *Save*.

7.2 Memory Usage

Monitor system performance, analyze workloads and memory history, and view load unit configuration.

The *Memory Usage* card lets you switch between a summary of used and resident memory. When you move your mouse over the summary graph, a category breakdown appears. If the database experiences out of buffer or out of memory events within the last 24 hours, a warning message indicating the number of each event appears below the *Memory Usage* card title.

7.2.1 Memory Analysis

Analyzing the memory allocation of the SAP HANA database can help you understand and resolve unusual memory usage and out-of-memory incidents.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

The *Memory Analysis* app enables you to visualize and explore the memory allocation of every service of a selected host during a specified time range. If you notice an increase in overall memory usage, you can investigate whether it's due to a particular component, allocator, or table.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click the *Memory Usage* *More (...)* then click *Analyze Memory History*.
2. Use the elements in the header to configure the display of memory statistics:

Dropdown	Description
Host and Services Selection	Opt to display memory statistics from a different host and service combination.
Unit Selection	Select a unit of measure in which to display all memory statistics

Dropdown	Description
Last selected server time	Select a time range for the memory allocation (upper) section, or select *Custom* in order to edit the time range.
(Navigate Application)	Opt to view the service you've selected through the Memory Analysis app in the Performance Monitor or the Workload Analyzer.
(Collection Frequency)	View how often information is collected to populate the various charts and graphs. If it's necessary, you can use the *Configuration of System Properties* app to modify the interval parameter of a statisticsserver component. Use caution because changes to the interval can affect performance.
(Memory Alert Settings)	Turn on or off the display of alerts and specify which priority level of alerts to display.
Threshold of Memory Usage	View the threshold of high priority memory usage alerts so you can determine what percentage of the effective allocation limit is being used by a service.

> **i Note**
>
> Data is collected and displayed only after the statistics server has been enabled. Ensure that you've configured the system properties (the .ini file) so that the statistics server's `active` parameter is set to `true`.

3. Analyze the memory statistics by exploring the data in the upper chart.

Element	Description
Allocated Memory	The pool of memory preallocated by the host for storing in-memory table data, thread stacks, temporary results, and other system data structures.
Host Allocation Limit	The `global_allocation_limit` for the host (as set in the `global.ini` configuration file).
Service Allocation Limit	Each service running on the host has an allocation limit. Collectively, all services can't consume more memory than the global allocation limit.
Total Used Memory	The total amount of memory used by the selected service on the selected host, including program code and stack, all data and system tables, and the memory required for temporary computations.
Memory Alert Icon	Represents one or more alerts triggered by a memory event. You can click individual alert icons and scroll through details, or navigate to the *Alerts* app.

4. Move the vertical selection bar in the upper chart to populate the data in the lower chart. The vertical selection bar snaps to the closest time for which there's collected data for the components. You can control what is displayed in the lower chart by selecting one of its top tabs, as described in the next steps.

5. Select the *Top Consumers* tab to have the lower chart display details about what is consuming the most memory for the host and service, in the given time period. Up to 50 top consumers are displayed. You can step through the collected data points by using the arrow buttons.

Element	Description
Name	Up to 50 allocators, objects, or other elements that have consumed the most memory during the specific time (chosen by the vertical selection bar in the upper chart).
Used Memory	The amount of memory consumed, displayed in the measurement units specified at the top of the page.
Component	The component with which the consumer is associated.
Top Consumers History	Expand *Top Consumers History* to see a chart of the top consumers (up to 10) for the time range selected at the top of the page.

6. Select the *Components* tab to have the lower chart display the *Used Memory by Component*.

Element	Description
Used Memory by Component	For the specific time (chosen by the vertical selection bar in the upper chart), the components of the selected service are listed in descending order of used memory.
Used Memory by Type	The donut chart displays a visual representation of the types of used memory for the specific time.
Components Used Memory History	Filling the checkbox of one or more components populates the *Used Memory History* chart.

7. Select the *Allocators* tab to display more detailed memory use in the lower chart. You can filter by component type. You can step through the collected data points by using the arrow buttons. (Allocators that have used less than 1 GB of memory aren't displayed.)

Element	Description
Used Memory by Allocator	For the specific time (chosen by the vertical selection bar in the upper chart), allocators of the selected component are listed in descending order of used inclusive memory. By clicking on an allocator, you can expand the list.
Filter by component name	To further refine the displayed allocator data, select the filter icon to specify one or more component names.
Allocators Used Memory History	Filling the checkbox of one or more allocators populates the Used Memory History chart.

8. Select the *Tables* tab to see statistics about tables' used memory or number of rows in the lower chart. Filter the display using the *Chart Value*, *Show Top (number)*, *Filter by Schema*, and *Time Range* filter options.

Element	Description
Top Tables by Size	When you select the *Used Memory Size* chart value, the chart displays the breakdown of memory usage of the highest consuming tables for the specified time.
Top Tables by Growth	When you select the *Used Memory Size* chart value, the chart displays the memory usage of the tables with the largest change in consumption for the selected time period. Hover over the data to see the *Previous Size* memory usage value from the beginning of the time period and the *Growth* during the time period (where the current size of the table is the sum of *Previous Size* and *Growth*).
Top Tables by Rows	When you select the *Number of Rows* chart value, the chart displays the number of rows for each of the highest consuming tables for the specific time (chosen by the vertical selection bar in the upper chart).
Top Tables by Rows Growth	When you select the *Number of Rows* chart value, the chart displays the number or rows for each of the tables with the largest change in consumption for the selected time period.

9. Select the *Out of Memory Events* tab to have the lower chart display the number of unique out-of-memory events that have occurred in the time range specified in the header. (The vertical selection bar doesn't influence the number of events displayed.)

Element	Description
Occurrences	The number of times a specific OOM event has been triggered.
Last Occurrence	The time and date of the most recent occurrence of the OOM event.
Last Reason	The parameter that triggered the most recent occurrence of the OOM event.
Statement	The SQL statement related to the OOM event.
Statement Hash	The unique identifier for the OOM event. Click the OOM identifier to open the Workload Analyzer and investigate the event.

> → Tip
>
> If an event has a corresponding OOM dump file, you can select *View Trace* to launch the Dump Viewer in the SAP Database Explorer.

In *Memory Statistics* charts, you can choose to display historical data for a time range between 24 hours and six weeks. In order to have a date range longer than six weeks (42 days), you can use SQL to update the `RETENTION_DAYS_CURRENT` value in the table `"_SYS_STATISTICS"."STATISTICS_SCHEDULE"`.

Related Information

Types of Memory Alerts [page 143]
System Views Used to Create Memory Alerts [page 146]
Memory Usage in the SAP HANA Database
Services [page 159]
Database Configuration in SAP HANA Cockpit [page 443]

7.2.1.1 Types of Memory Alerts

The *Memory Statistics* chart can notify you about a variety of memory-related alerts.

Memory alerts related to the entire system

Alert	Description	Threshold Units	ALERT_ID
Licensed memory usage	Determines what percentage of licensed memory is used.	percent	44
Overflow of metadata version space	Determines the overflow ratio of the metadata version space.	ratio	74
Agent memory usage	Determines what percentage of total memory available on the agent is used.	percent	701

Memory alerts related to the host

Alert	Description	Threshold Units	ALERT_ID
Host physical memory usage	Determines what percentage of total physical memory available on the host is used. All processes that consume memory are considered, including non-SAP HANA processes.	percent	1
Streaming project physical memory usage	Determines what percentage of total physical memory available on the host is used for the streaming project.	percent	602
Plan cache hit ratio	Determines whether the plan cache hit ratio is too low.	ratio	91
Cached view size	Determines how much memory is occupied by the cached view.	percent	81

Memory alerts related to services

Alert	Description	Threshold Units	ALERT_ID
Memory usage of name server	Determines what percentage of allocated shared memory is being used by the name server on a host.	percent	12
Memory usage of services	Determines what percentage of its effective allocation limit is being used by a service.	percent	43
Plan cache hit ratio	Determines whether the plan cache hit ratio is too low.	ratio	91
Cached view size	Determines how much memory is occupied by the cached view.	percent	81

Memory alerts related to tables

Alert	Description	Threshold Units	ALERT_ID
Record count of non-partitioned column-store tables	Determines the number of records in non-partitioned column-store tables, but ignores information stored in statistics_exclude_tables. Current table size isn't critical. Partitions need only be considered if tables are expected to grow rapidly. (A non-partitioned table can't contain more than 2,000,000,000 (2 billion) rows).	records	17
Alert_Mon_Column_Tables_Record_Count_Incl	Determines the number of records in non-partitioned column-store statistics_exclude_tables. This alert is inactive, unless you activate it.	records	117
Table growth of non-partitioned column-store tables	Determines the growth rate of non-partitioned column tables.	percent	20
Record count of column-store table partitions	Determines the number of records in the partitions of column-store tables, but ignores information stored in statistics_exclude_tables. A table partition can't contain more than 2,147,483,648 (2 billion) rows.	records	27

Alert	Description	Threshold Units	ALERT_ID
Alert_Partitioned_Table _Record_Count_Incl	Determines the number of records in the partitions of column-store statistics_exclude_tables. This alert is inactive, unless you activate it.	records	127
Size of delta storage of column-store tables	Determines the size of the delta storage of column tables.	MB	29
Total memory usage of column-store tables	Determines what percentage of the effective allocation limit is being consumed by individual column-store tables as a whole (that is, the cumulative size of all of a table's columns and internal structures).	percent	40
Memory usage of main storage of column-store tables	Determines what percentage of the effective allocation limit is being consumed by the main storage of individual column-store tables.	percent	45
Columnstore unloads	Determines how many columns in columnstore tables have been unloaded from memory. This alert can indicate performance issues.	tables	55
Total memory usage of table-based audit log	Determines what percentage of the effective memory allocation limit is being consumed by the database table used for table-based audit logging. If this table grows too large, the availability of the database could be impacted.	percent	64
Table growth of rowstore tables	Determines the growth rate of rowstore tables.	percent	67
Total memory usage of row store	Determines the current memory size of a row store used by a service.	percent	68
Row store fragmentation	Check for fragmentation of row store.		71
Overflow of rowstore version space	Determines the overflow ratio of the row store version space.	ratio	73
Rowstore version space skew	Determines whether the row store version chain is too long.	version	75

Alert	Description	Threshold Units	ALERT_ID
Auto merge of column-store tables	Determines whether the delta merge of a table was executed successfully.	records	88

7.2.1.2 System Views Used to Create Memory Alerts

The SAP HANA cockpit displays memory alerts in the *Memory Analysis* app based on SAP HANA system views.

The following system views provide the information with which values for current and historical memory allocation are calculated:

Views from the _SYS schema:

- M_HOST_RESOURCE_UTILIZATION
- M_SERVICE_MEMORY
- M_SERVICE_COMPONENT_MEMORY
- M_RS_TABLES
- M_HEAP_MEMORY
- M_CS_COLUMNS
- M_OUT_OF_MEMORY_EVENTS
- M_SQL_PLAN_STATISTICS

Views from the _SYS_STATISTICS schema:

- HOST_RESOURCE_UTILIZATION_STATISTICS
- HOST_SERVICE_MEMORY
- HOST_SERVICE_COMPONENT_MEMORY
- HOST_HEAP_ALLOCATORS
- GLOBAL_ROWSTORE_TABLES_SIZE_BASE
- HOST_COLUMN_TABLES_PART_SIZE
- STATISTICS_OBJECTS
- STATISTICS_CURRENT_ALERTS
- STATISTICS_ALERT_INFORMATION
- STATISTICS_ALERT_THRESHOLDS
- GLOBAL_OUT_OF_MEMORY_EVENTS

For more information about these views, see the *SAP HANA SQL Reference Guide*.

Related Information

SAP HANA SQL Reference Guide for SAP HANA Platform

7.2.2 Managing Warm Data With the Native Storage Extension

SAP HANA native storage extension (NSE) is a general-purpose, built-in warm data store in SAP HANA that lets you manage less-frequently accessed data without fully loading it into memory.

SAP HANA NSE integrates disk-based database technology (for example, disks and SSDs) with the SAP HANA in-memory database for an improved cost-to-performance ratio, while complementing other warm data tiering solutions such as SAP HANA Extension Node and SAP HANA dynamic tiering.

For more information, see *SAP HANA Native Storage Extension* in the *SAP HANA Administration Guide*.

Related Information

SAP HANA Native Storage Extension

7.2.2.1 Monitor Memory Paging for Non-Partitioned Tables

Provides paging information (in-memory, in buffer cache, unloaded) at column level for non-partitioned tables.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Use the *Memory Paging Monitor* to view information about where data in your database tables is being stored.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected,click the ••• *(More)* icon on the *Memory Usage* card, and then *Monitor Memory Paging*.

   SAP HANA Cockpit displays a list of partitioned and non-partitioned tables. By default, the top 10 tables in the database are shown.
2. Click on the required non-partitioned table either in the chart or the table section.

For each column, the *Column Sizes* bar chart displays the column data that is in-memory, in buffer cache, or unloaded.

Results

If you want to change how your data is stored, click *View Load Units Configuration* to view and alter the load unit configuration for your tables, partitions, or columns.

7.2.2.2 Monitor Memory Paging for Partitioned Tables

Review memory paging statistics for all partitions and columns of a selected partitioned table in SAP HANA native storage extension.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected,click the ••• *(More)* icon on the *Memory Usage* card, and then *Monitor Memory Paging*.

   SAP HANA Cockpit displays a list of partitioned and non-partitioned tables.
2. Click a partitioned table in the *Table* list.

   You now have the option of viewing memory statistics for either the table's partitions or the table's columns.
3. Choose one of the following:

Option	Description
Click *Partitions*	In the *Partitions* chart, view information about the size and paging status of the table's partitions. Click a partition to view the list of columns for the selected partition.

Option	Description
Click *Columns*	In the *Columns* chart, view information about the size and paging status for each column. Click a column to view information on each column, such as total memory size, on-disk size, in-memory and buffer cache sizes, and load unit.

Results

If you want to change how your data is stored, click *View Load Units Configuration* to view and alter the load unit configuration for your tables, partitions, or columns.

7.2.2.3 View and Modify Load Unit Configuration

View information on configured load units for SAP HANA NSE tables, partitions, sub-partitions, and columns and modify load unit configurations using SAP HANA Cockpit.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Loading behavior is determined by the load unit (one of PAGE, COLUMN, and DEFAULT) specified for the column, partition, and table in SAP HANA native storage extension (NSE) column-store tables.

A column can be an in-memory column (which has a load unit of COLUMN) or a paged-attribute column (which has a load unit of PAGE), depending on its specified load unit. When you run a DDL (for example, ALTER TABLE), SAP HANA determines the effective load unit for a column based on load unit precedence.

To determine whether a column is fully in-memory or is paged, SAP HANA checks the load unit that is being set at the column level.

> **i** Note
>
> DEFAULT indicates that the user doesn't have any explicit load unit preference for a database object. A database object with the tag as DEFAULT LOADBALE inherits the load unit preference of its parent object. For example, a column inherits its partition's preference and a partition inherits its table's load unit preference. The system default load unit is COLUMN LOADABLE.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click the ••• *(More)* icon on the *Memory Usage* card, and then *View Load Unit Configuration*.

   The *Load Unit Configuration* page opens, showing the following information (filterable by schema and/or table):

Column	Description
Part ID	The partition identifier if the table is partitioned.
Range	The range type of the partitioned table. Can be one of RANGE, or RANGE-RANGE.
Effective Unit	The effective load unit for a column based on load unit precedence and can be one of ALL COLUMN, MIXED, COLUMN, or UNLOADED. If a column's effective unit is displayed as UNLOADED, it indicates that the table is either not loaded or a non-logging table.
Load Unit Setting	The unit that determines the loading behavior. Can be one of PAGE, COLUMN, or DEFAULT. If you have the DATA ADMIN privilege, then you can change the load unit setting by clicking the setting, selecting the type of load unit you want to change it to and clicking *Save*.

2. (Optional) Execute SQL statements to update the load units.
   a. Click *Open SQL Console* on the *Load Unit Configuration* page to open the SAP HANA database explorer.
   b. Execute one or more of the following statements:

Statement	Description
`ALTER TABLE … ALTER LOAD UNIT`	Alters the load unit configuration for an entire table. This command affects all partitions and columns in the table.
`ALTER TABLE … ALTER LOAD UNIT`	Alters the load unit configuration for one or more table partitions.
`ALTER TABLE … ALTER … <alter load unit column syntax>`	Alters the load unit configuration for one or more table columns.

Related Information

Changing the Load Unit Using ALTER TABLE

7.2.2.4 View and Configure the Native Storage Extension Buffer Cache Monitor

Use the *Buffer Cache Monitor* to view memory usage statistics for the native storage extension buffer cache and determine which tables, partitions, and columns are using most of the cache.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

The *Buffer Cache Monitor* is a useful tool for analyzing if you've received an out-of-buffer alert. Out-of-buffer alerts appear immediately below the title on the Memory Usage card. You can also view the size of a used cache versus its configured size.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected,click the ••• *(More)* icon on the *Memory Usage* card, and then *Monitor Buffer Cache*.

   The *Buffer Cache Size* chart displays used and configured buffer cache sizes. For each buffer cache entry, the following information is provided:

Column	Description
Host/Port	The host name and internal port of the HANA service.
Server	The indexserver of each database on the SAP HANA server.
Volume ID	The volume ID for the HANA service. Volume IDs are unique to a HANA service and it uniquely identifies the data and log volumes for its mapped service. The mapping between the volume ID and the SAP HANA service is listed in the M_SERVICES system view.
Cache Name	The system-generated name for the buffer cache.
Warning	Displays any warnings for the buffer cache. For example, the active buffer cache has been disabled but is still being accessed by users.

Column	Description
State	The current state (enabled or disabled) of the buffer cache.
Replacement Policy	The caching algorithm (IMPROVED LRU) to manage memory pages in buffer cache.
Out of Buffers Count	The number of times the NSE buffer cache wasn't large enough to handle the workload causing SAP HANA to generate an out-of-buffer alert. During an out-of-buffer situation, regular user tasks can be rolled back, and critical tasks use the emergency buffer pool.
Configured Cache Size	The user configured buffer cache size. The buffer cache is enabled by default and its size is 10% of the SAP HANA memory.
% Used	The percentage of cache that has been used.
Buffer Reuse Count	The number of times the buffers for used pages have been recycled to provide memory for new pages; the time is calculated since the last server restart.
Hit Ratio	The ratio of the number of cache hits to the number of lookups. This value indicates the success rate of pages found in buffer cache.

2. Clicking on a specific buffer cache row in the *Buffer Caches* section opens the *Memory Paging Monitor* page and shows a list of tables for the host that are associated with the selected cache.
3. Click *Configure Buffer Cache Size* to configure the cache size by either setting a maximum cache size as a percentage of total SAP HANA memory, setting a maximum cache size cap (in MB), or setting both. Specifying values for these settings alters the relevant buffer cache size parameter values in the corresponding `.ini` file.

Related Information

SAP HANA Native Storage Extension Buffer Cache

7.2.2.5 Monitor Out of Buffer Events

View a list all out of buffer events that occurred during the specified time range.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

On the *Database Overview* page, with the *Monitoring* or *All* view selected, click the ... *(More)* icon on the *Memory Usage* card, and then *View All Out of Buffer Events*.

> **i** Note
>
> The menu option only appears if out of buffer events have occurred.

The *Memory Analysis* chart displays out of buffer events. For each out of buffer entry, the following information is provided:

Column	Description
Host/Port	Specifies the host name and internal port of the HANA service.
Type	Specifies the type of the event.
Reason	Specifies additional information for the event (for example, the number of buffer cache used and how many buffer cache are needed).
Creation Time	Specifies the time the event occurred.
Update Time	Specifies the time the event was last updated.
Handle Time	Specifies the time a user performed an operation that resolved or alleviated the issue related to the event.
State	Specifies the current state of the event (for example, NEW, HANADLED).

7.2.3 View and Configure the Persistent Memory Monitor

Use the Persistence Memory Monitor to configure persistent memory in your SAP HANA database, and to monitor and manage data stored in the persistent memory.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

You must have set up the persistent memory feature during the installation of your SAP HANA database by using the SAP HANA database lifecycle manager (HDBLCM). Using HDBLCM you must set the `use_pmem` parameter to enable persistent memory and then set the `pmempath` parameter to set the `basepath`. For more information on the lifecycle manager, refer to the *SAP HANA Server Installation and Update Guide*.

Context

Persistent memory (PMEM), also known as storage-class memory, is byte-addressable and is similar to DRAM in its latency characteristics and also in that it's non-volatile. PMEM can be treated by the CPU as RAM, offering fast read and write performance.

Use persistence memory to achieve ultra-fast access to large datasets and to ensure that data persists in memory even after power interruption.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Memory Usage* > *Configure Persistent Memory*.

   You can also configure persistent memory by clicking the *Configure Persistent Memory* button on the Persistent Memory Monitor.

   a. Click *Add* to turn on persistent memory and to specify the location of the persistent memory file system.

   > **i Note**
   >
   > Specifying a file system path overrides the `table_default` parameter in the `persistent_memory` section of the `indexserver.ini` file and the `basepath_persistent_memory_volumes` parameter in the `persistence` section of the `global.ini` file. Restart your system after altering these parameters for the new settings to take effect.

2. Monitor persistence memory usage.
3. The *Table Sizes* view in the Persistent Memory Monitor shows the bytes stored in persistent memory for the top 10, 20, or 50 user tables. You can also choose to view the persistent memory used per schema, host, or schema and host.
4. To view the column layer for a non-partitioned table that is stored in persistent memory, click the table name in either the chart or the list.

   In the *Column Sizes* view, you can also see the size of the column, and the name of the file that saves the persistent data for each column.
5. If the table is partitioned and you click on the table name, then you see the table's persistent memory usage in the partition layer as well as the partition ID.

   You can also access the *Column Sizes* view for a partitioned table to see the partition's persistent memory usage for each column.

Related Information

SAP HANA Server Installation and Update Guide

7.2.3.1 Persistent Memory

Persistent memory (non-volatile RAM, also referred to as Storage Class Memory) is supported in SAP HANA as a persistent storage type.

Persistent memory (or NVRAM) is an emerging class of memory that combines the qualities of both DRAM and Flash storage and bridges the gap between disk storage and main memory. For the purposes of SAP HANA, its most important characteristics are that it's byte addressable like DRAM and can be treated by the CPU as RAM, offering fast read and write performance; its latency characteristics are also very close to DRAM.

The latest information about this feature, such as hardware availability and operating system support, is provided in SAP Note 2618154 *SAP HANA Persistent Memory (NVM)*.

The persistent storage functions as DRAM memory and is used specifically for MAIN data fragments of the column store where approximately 90% of the data of a table is held. In this way, DELTA data can continue to be stored in DRAM (a relatively small quantity of data requiring the fastest possible access) with the MAIN segment of a table in NVRAM. Note that if persistent memory storage is enabled, data is still written to the data volumes but is ignored. After a system restart, the main data fragments saved in persistent storage are still available in memory and don't have to be reloaded. This storage type has no impact on backup and recovery or high availability processes.

As with the data volume persistent storage feature, data is persisted on the basis of savepoints. In the case of non-volatile memory storage, NVM blocks are constructed whenever the in-memory MAIN fragment is created and are kept consistent according to HANA's savepoints.

The key performance benefit is in greatly accelerated start-up times so that data can be quickly reloaded into memory and any system downtime can be minimized. Other benefits can include higher memory capacities per server.

Performance differences of this storage type compared to DRAM due to higher latency are mitigated by features such as hardware prefetchers. It's also possible to enable persistent memory storage for selected tables or for specific columns of a table so that a selective implementation of this storage type is possible (see the Configuration section).

Installation

The persistent memory feature can be set up during system installation using the SAP HANA database lifecycle manager (HDBLCM). The lifecycle manager uses two parameters to enable persistent memory (`use_pmem`) and the `pmempath` parameter to set the basepath (see the Configuration section). For more information on lifecycle manager, refer to the *SAP HANA Server Installation and Update Guide*.

At the level of the Linux operating system, SAP HANA only supports DAX enabled (Direct Access) file systems.

Filesystem layout of persistent memory block storage

The directory specified as the basepath (see Configuration) is the root persistent memory storage directory beneath which data for each savepoint is stored. The structure beneath the basepath is maintained by HANA and no attempt should be made to modify this content. Multiple logical versions of blocks are saved and are organized in subdirectories named on the basis of, for example, service name, Volume ID and current savepoint version. There are two subdirectory systems, one for data and one for deletion requests: the memory storage persistence maintains a history of requests to delete blocks in a *tombstones* directory parallel to the *data* directory. In the case of tombstone blocks, the savepoint version relates to the savepoint when the data block was requested for deletion.

Configuration

The default behavior for using persistent memory is determined automatically: the system tests for the DAX filesystem at the defined persistent memory basepath. If the hardware is installed and the basepath correctly configured pointing to the persistent memory mount, then alltables use persistent memory by default. The SAP HANA installer sets the persistent memory basepath parameter if it's able to identify the DAX filesystem and if the user of the installation tool confirms the usage; the basepath parameter can also be set manually if necessary.

This default behavior can be overridden at four levels in the sequence shown here:

4	Database	Can be enabled or disabled by configuration parameter.
3	Table	Can be enabled or disabled by SQL statement.
2	Partition	Can be enabled or disabled by SQL statement.
1	Column	Can be enabled or disabled by SQL statement.

This level of granularity offers a high degree of flexibility: if persistent memory is applied at a certain level it's automatically inherited at lower levels of the sequence but can also be overridden at lower levels.

At the highest configuration level (database), persistent memory is managed by setting the `table_default` configuration parameter for the database as a whole. You can switch `table_default` to OFF to enable more selective persistent memory storage at lower levels for particular tables and partitions.

At other levels, persistent memory is managed at the SQL command line using the `alter_persistent_memory_spec` clause with the CREATE / ALTER table commands. Here, we give examples to show how it can be applied at each level; refer to the SAP HANA SQL Reference Guide for full details of this clause.

Configuration Parameter: Basepath

As with the other data and log persistent storage volumes and also the fast restart option (tmpfs), the storage location basepath for persistent memory storage must be defined as a configuration parameter in the `persistence` section of the global.ini file. Enter the basepath location in the following parameter:

Parameter	**basepath_persistent_memory_volumes**
Short Description	Location of NVRAM storage.
Full Description	Data for NVRAM-enabled tables or columns is stored at the location defined in this parameter. Multiple locations can be defined here using a semicolon as a separator (no spaces). Multiple paths correspond to NUMA nodes without any ordering restriction in terms of NUMA node mapping, for example, for a four-NUMA node system: **"**`/mnt/pmem0;/mnt/pmem1;/mnt/pmem4;/mnt/pmem2`
Type	Text
Change	Restart required (offline changeable only)
Default	Blank

Parameter	basepath_persistent_memory_volumes
Layer Restriction	HOST - can only be set at the HOST level.

Configuration Parameter: Enable persistent memory storage

The `table_default` parameter is in the `persistent_memory` section of the indexserver.ini file.

Parameter	table_default
Short Description	Enable or disable persistent memory storage for the database generally.
Full Description	This parameter can be set to ON or OFF or DEFAULT (no explicit preference, follow the existing default behavior). The setting applies for the complete database but can be overridden by settings applied at a lower level.
Change	Online
Default	DEFAULT

SQL Alter Persistent Memory Clause

Persistent memory can be enabled by SQL command using the `alter_persistent_memory_spec` clause with CREATE and ALTER table. This command can be used at any of the three levels (table, partition, or column) to enable or disable persistent memory storage using the PERSISTENT MEMORY switch. The following examples illustrate this method using CREATE TABLE.

Example 1: Create table with persistent memory storage enabled for the new table.

```
CREATE COLUMN TABLE PMTABLE (C1 INT, C2 VARCHAR (10)) PERSISTENT MEMORY ON
```

PERSISTENT MEMORY can be set to ON, OFF, or DEFAULT meaning no preference.

Example 2: Create range-partitioned table with persistent memory storage for a selected partition (only valid for range partitions):

```
 CREATE COLUMN TABLE PMTABLE (C1 INT) PARTITION BY RANGE (C1) (
                   PARTITION '0' <= VALUES < '10' PERSISTENT MEMORY ON,
                   PARTITION OTHERS PERSISTENT MEMORY OFF);
```

Example 3: Create table with persistent memory storage applied to selected columns of a table:

```
CREATE COLUMN TABLE PMTABLE (C1 INT PERSISTENT MEMORY ON, C2 VARCHAR (10), C3
INT PERSISTENT MEMORY OFF)
```

For ALTER TABLE, a preference value is also required (in addition to the ON/OFF switch) to determine how the change is applied. This clause is the `alter persistent memory preference` clause that requires one (or more) of the following keywords:

- IMMEDIATE - the change is applied immediately and the specified storage is populated.
- DEFERRED - the specified storage will be populated at the time of the next delta merge or reload.
- CASCADE - this keyword can be used at table level in addition to IMMEDIATE and DEFERRED to apply the change to all lower levels of the hierarchy.

The following example shows this usage with ALTER TABLE and cascades the change immediately to partitions and columns of the named table:

```
ALTER TABLE MYTABLE PERSISTENT MEMORY ON IMMEDIATE CASCADE
```

Unloading Tables from Memory

The standard UNLOAD statement, and also table unloads, executed implicitly by HANA services due to memory running low, don't by default free the memory but simply remove the mapping of relevant data structures from the HANA service process address space and the MAIN data fragment remains; this applies to both persistent memory and the fast restart option (tmpfs). The configuration parameter `table_unload_action` in the [persistent_memory] section of the indexserver.ini file can be used to control this unload behavior. The default value of the parameter is RETAIN but this can be set to DELETE in which case MAIN data fragments are cleared from memory.

Regardless of the setting of this parameter, you can also use UNLOAD with the additional DELETE PERSISTENT MEMORY clause to free memory, for example:

```
UNLOAD FRAOTABLE DELETE PERSISTENT MEMORY;
```

Monitoring Views

For full details of the monitoring views and values named here, refer to the *SAP HANA SQL and System Views Reference*.

The following monitoring views show information about persistent memory in the PERSISTENT_MEMORY column: M_CS_TABLES, M_CS_COLUMNS and M_CS_ALL_COLUMNS. This information shows either TRUE (persistent memory is enabled) or FALSE.

At the level of tables (M_CS_TABLES), usage information is shown as PERSISTENT_MEMORY_SIZE_IN _TOTAL.

At the level of columns (M_CS_COLUMNS and M_CS_ALL_COLUMNS) usage information is shown in a series of columns named: MAIN_PERSISTENT_MEMORY_SIZE_IN_*. These columns show usage in the categories DATA, DICT, INDEX, MISC, and UNUSED. These column views also include STORED_IN_PERSISTENT_MEMORY and LOADED_FROM_PERSISTENT_MEMORY, both TRUE/FALSE values.

Three monitoring views show full statistical details of persistent memory usage:

- M_PERSISTENT_MEMORY_VOLUMES - capacity, usage, and metadata of persistent memory volumes configured per NUMA Node.
- M_PERSISTENT_MEMORY_VOLUME_DATA_FILES - metadata statistics about files created by SAP HANA services for data storage on the persistent memory volumes.
- M_PERSISTENT_MEMORY_VOLUME_STATISTICS - statistics of physical lifecycle events of blocks managed by SAP HANA services on the persistent memory volumes.

The view M_TABLE_PERSISTENT_MEMORY_FILES shows persistent memory file information for all of the tables, partitions, and columns in the database.

Statistics server collectors save historic information in the following statistics views:

- _SYS_STATISTICS.HOST_PERSISTENT_MEMORY_VOLUMES - reports the capacity, usage, and metadata of persistent memory volumes.
- _SYS_STATISTICS.HOST_PERSISTENT_MEMORY_VOLUME_STATISTICS - reports the statistics of physical lifecycle events of blocks managed by SAP HANA services on the persistent memory volumes.

Related Information

SAP Note 2618154
SAP HANA Server Installation and Update Guide
SAP HANA Fast Restart Option

7.3 Services

To monitor the health of your SAP HANA database in more detail, for example, to troubleshoot performance bottlenecks, you can analyze the status and resource usage of individual database services. If necessary, you can perform follow-up operations, such as starting missing services, stopping a service, or killing a service. You can also start or stop a system.

You can use the cockpit to monitor and manage more than one database, each running version SAP HANA 1.0 SPS 12 or later. Any database running version SAP HANA 2.0 SPS 01 or later is set in multiple-container mode by default. The cockpit can also monitor single-container systems running earlier versions of SAP HANA. When you access the *Manage Services* application, the operations you have the option to perform depend on whether you are displaying a tenant or a system database.

Click *Manage Services* on the *Services* card. The *Services* card appears on the *Database Overview* page when the *Monitoring* or *All* view is selected.

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*.

Not all columns are visible by default. You can configure which columns are visible by clicking the ⚙ *(Settings)* icon. You can configure the sort order of the information by clicking the ↑↓ *(Sort)*icon.

If there are any alerts in the system, you can see details by clicking the alert in the *Service Alerts* column or by clicking *Go to Alerts* at the top of the page.

If you want to investigate the memory usage history of a particular service, click the mini chart in the *Memory* column to open *Memory Analysis* for the service. You can also use *Reset Memory Statistics* link to clear the statistics history.

Depending on the situation, you may need to perform further operations on all or selected services (for example, start, stop, or kill a service). For more information about the available options, see *Operations on Services*.

If necessary, you can also start or stop a system. See *Start a Database* and *Stop a Database*.

Related Information

Alerts [page 126]
Memory Analysis [page 139]
Start a Database [page 164]
Stop a Database [page 166]
Add or Remove Services in a Tenant Database [page 422]

7.3.1 Service Details

Manage Services provides you with detailed information about database services for an individual database.

> **i Note**
>
> Not all of the columns listed below are visible by default. You can add and remove columns in the Columns dialog, which you open by clicking the ⚙ *(Settings)* icon in the table toolbar.

The following table lists the information available for services.

Column	Description
Host	Name of the host on which the service is running
Service	Service name, for example, indexserver, nameserver, xsengine, and so on
Status	The status of the service The following statuses are possible: • *Running* • *Running with Issues* (where at least one service isn't running, or there is at least one high alert) • *Starting* • *Stopping* • *Stopped* • *Not Running* To investigate why the service isn't running, you can navigate to the crash dump file created when the service stopped. **i Note** The crash dump file opens in the *Trace* tool of the SAP HANA Web-based Development Workbench. To see this information, you need the role `sap.hana.xs.ide.roles::TraceViewer` or the parent role `sap.hana.xs.ide.roles::Developer`.

Column	Description
Role	Role of the service in a failover situation Automatic failover takes place when the service or the host on which the service is running fails. The following values are possible: • *Master* The service is the active master worker. • No entry The service is a slave worker. • *Standby* The service is in standby mode. It doesn't contain any data and doesn't receive any requests.
Port	Port that the system uses for internal communication between services
Start Time	Time at which the service started **i Note** The time is given in the timezone of the SAP HANA server.
Service Alerts	Alerts triggered for the service.
Process ID	Process ID
CPU	Mini chart visualizing the CPU usage of the service Clicking the mini chart opens the *Performance Monitor* for a more detailed breakdown of CPU usage.
Memory	Mini chart visualizing the memory usage of the service • Dark green shows the service's used memory. • Light green shows the service's peak memory. • The gray stroke represents the effective allocation limit. • The light gray background represents the physical memory. Clicking the mini chart opens the *Memory Analysis* app for a more detailed breakdown of memory usage.
Used Memory (MB)	Amount of memory currently used by the service Clicking the mini chart opens the *Memory Analysis* app for a more detailed breakdown of memory usage.
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
Physical Memory 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

Column	Description
Allocated Shared Memory (MB)	Shared memory part of the allocated memory pool
Allocation Limit (MB)	Maximum size of allocated memory pool
CPU Process Usage (%)	CPU usage of process
CPU Host (%)	CPU usage on host
Virtual Memory 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
Action	The available action for the service
Reset Memory Statistics	Resets all memory statistics for all services. Peak used memory is the highest recorded value for used memory since the last time the memory statistics were reset. This value is useful for understanding the behavior of used memory over time and under peak loads. Resetting peak used memory allows you, for example, to establish the impact of a certain workload on memory usage. If you reset peak used memory and run the workload, then you can then examine the new peak used memory value.
Go To Alerts	Displays the alerts for this database.

Related Information

Memory Usage in the SAP HANA Database
Memory Analysis [page 139]
The Performance Monitor [page 515]
Monitoring Alerts

7.3.2 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.

Some tenant services, such as adding and starting missing services, can only be managed when connected to the SYSTEMDB, while other tenant services, such as killing a service or removing a service from a tenant can be managed when connected to either the SYSTEMDB or the tenant.

> **i Note**
>
> - Depending on the service, some options may not be available. You can use the cockpit to monitor and manage more than one database, each running version SAP HANA 1.0 SPS 12 or later. Any database running version SAP HANA 2.0 SPS 01 or later is set in multiple-container mode by default. The cockpit can also monitor single-container systems running earlier versions of SAP HANA.
> - 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*.

Start and Stop Services in a Tenant Database

Option	Description	Connected As	See
Start Missing Services	Starts any inactive services.	SYSTEMDB only.	Start and Stop Services in a Tenant Database
Stop Service	Stops the selected service normally. The service then typically restarts automatically.	SYSTEMDB or tenant.	Start and Stop Services in a Tenant Database
Kill Service	Stops the selected service immediately, and if *Generate crash dump file* is selected, creates a crash dump file. The service then typically restarts automatically.	SYSTEMDB or tenant.	Start and Stop Services in a Tenant Database
Add Service	Adds the service you select from the list. Services can only be added to a tenant, not the system database itself.	SYSTEMDB only.	Add or Remove Services in a Tenant Database

Option	Description	Connected As	See
Remove Service	Removes the selected service. You can only remove services that have their own persistence. If data is still stored in the service's persistence, it's redistributed to other services. You can't remove the following services: • Master name server • Master index server • Primary index server on a host **i Note** To remove a service, you must have the EXECUTE privilege on the stored procedure SYS. UPDATE_LANDSCAPE_CONFIGURATION.	SYSTEMDB or tenant.	Add or Remove Services in a Tenant Database

Related Information

Memory Usage in the SAP HANA Database
High Availability for SAP HANA
Start a Database [page 164]
Stop a Database [page 166]
Work with Alerts
Generate and Execute a Table Redistribution Plan [page 506]

7.3.3 Start a Database

Use SAP HANA cockpit to start a database.

Prerequisites

- You are connected to the SYSTEMDB.
- You have the credentials of the operating system user (`<sid>adm` user) that was created when the system was installed.

Procedure

On the *Database Directory*, click the *Start System* button at the top of the page.

Results

The database services start one by one, including those of any tenant databases. For details on starting an individual tenant database, see *Start a Tenant Database*.

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

> → Tip
>
> To analyze any problems that may occur during startup, you can access the system's diagnosis files from the homepage of the SAP HANA cockpit.

If you're unable to start a database that was registered while it was unreachable, check the information entered during registration. The cockpit can't check the registration information for an unreachable database, and thus can't tell the difference between a host or database that's unreachable and one that doesn't exist. In particular, make sure these are correct:

- Host name
- Instance number
- Technical user name and password
- SAP HANA system ID

If you find an error, unregister the database and register it again.

Related Information

Start a Tenant Database [page 383]
Services [page 159]
Register a Database [page 48]
Unregister a Database [page 56]

7.3.4 Stop a Database

Use SAP HANA cockpit to stop a database.

Prerequisites

- You are connected to the SYSTEMDB.
- You have the credentials of the operating system user (`<sid>adm` user) that was created when the system was installed.

Procedure

1. On the *Database Directory*, click the *Stop System* button at the top of the page.
2. Specify how you want to stop the system:

Option	Description
Softly	The system is stopped after all running statements have finished. If the system doesn't stop before the specified timeout, it is stopped immediately. The default timeout is 5 minutes.
Immediately	The system is stopped immediately. Open transactions are aborted and rolled back.

Results

The database services stop one by one. The services of tenant databases are stopped. For more information about how to stop an individual tenant database, see *Stop a Tenant Database* in the *SAP HANA Administration Guide*.

When all services have stopped, the system has the status *Stopped*.

> → Tip
>
> To analyze problems even when the system is stopped, you can access the system's diagnosis files from the homepage of the SAP HANA cockpit.

Related Information

Services [page 159]
Stop a Tenant Database [page 384]

7.4 System Replication

Determine whether a database is part of a system replication configuration, and monitor the status of replication between the primary system and the secondary systems.

7.4.1 Monitoring SAP HANA System Replication

To monitor SAP HANA system replication, you can use the *System Replication* card in the SAP HANA cockpit.

To open the *System Replication Overview* page, with the *Monitoring* or *All* view selected, click the *System Replication* card on the *Database Overview* page in the SAP HANA cockpit.

The *System Replication Overview* displays a graphical representation of the system replication landscape with the following information:

- The name and role of the system, as well as the selected operation mode
  For the operation modes `logreplay` and `logreplay_readaccess`, a retention time estimation is also displayed. The *Estimated log retention time* is an estimation of the time left before the primary system starts to overwrite the *RetainedFree* marked log segments and a full data shipping becomes necessary to get the primary and secondary systems back in sync after a disconnect situation. The *Estimated log full time* is an estimation of the time left before the primary system runs into a log full. The value shown in the header shows the situation into which the system could run first: log retention or log full.
- If the SQL ports of the secondary system are open for read access
- The replication mode used between the systems
- The current average redo log shipping time and the average size of shipped redo log buffers

Furthermore, detailed information on system replication is provided in the following tabs:

> **i Note**
>
> These tabs are displayed only if you configured a system replication before.

Tabs on the System Replication Overview

Tab Name	Description
Graphical Overview	The *Graphical Overview* tab provides information about each tier.
Related Alerts	The *Related Alerts* tab provides a description of any existing alerts, as well as their priority. This tab is only displayed when system replication-related alerts are available.
Replicated Services	The *Replicated Services* tab provides information on the replication status per site and service.

Tab Name	Description
Network	The *Network* tab provides information on the time it took to ship the redo log to the secondary system and to write the redo log to the local log volume on disk. You can select the network connection that you want to analyze (for example, *Network Site 1 to 2* or *Network Site 2 to 3*). The graph displayed compares the local write wait time with the remote write wait time monitored over the last 24 hours.
Log Shipping Backlog	The *Log Shipping Backlog* tab provides a graphical representation on the history of the log shipping backlog.
Log Replay	The *Log Replay* tab provides a graphical representation on the delay of the secondary system. This tab is displayed if the chosen operation mode for the system replication landscape is `logreplay` or `logreplay_readaccess`. When this tab is activated for a secondary system, the log replay delay is shown for the last 24 hours. Furthermore, in this tab you can select to visualize the estimated log retention time as well as the estimated log full time for all system replication relevant services.
Network Speed Check	The *Network Speed Check* tab provides a way to measure the network speed of the system replication host-to-host network channel mappings.
Time Travel	The *Time Travel* tab provides an overview of the available snapshots, as well as the possibility to start a secondary time travel. This tab is only displayed on the overview of the secondary system.
Takeover History	The *Takeover History* tab provides an overview of the entire takeover and time travel history available.
Parameter Settings	The *Network Security Settings* option displays the specific network security details configured between the primary and the secondary systems. The *Invisbile Takeover* option provides the possibility to turn on the invisible takeover. Differently from a standard takeover, an invisible takeover ensures that the client reconnects to the primary system and the sessions are restored to the secondary system.

7.4.1.1 SAP HANA System Replication Details

Detailed information from the M_SERVICE_REPLICATION and the M_SYSTEM_REPLICATION monitoring views about system replication.

General Overview

Column	Description
Site ID 1	Generated ID of the primary site
Secondary Site ID 2	Generated ID of the secondary site

Column	Description
Service	Name of the service
Volume ID	Persistence volume ID
Operation Mode	• LOGREPLAY • LOGREPLAY_READACCESS • DELTA DATA SHIPPING
Replication Mode	Configured replication mode: • SYNC: synchronous replication with acknowledgment when buffer has been written to disk on the secondary system • SYNCMEM: synchronous replication with acknowledgment when buffer arrived in memory on the secondary system • ASYNC: asynchronous replication where the primary doesn't wait for the acknowledgment • UNKNOWN: is set if replication mode can't be determined (for example, if there are communication errors when getting status information from a service).
Replication Status	Current status of replication: • UNKNOWN: secondary didn't connect to primary since last restart of the primary • INITIALIZING: initial data transfer is running. In this state, the secondary is first usable when this is finished • SYNCING: secondary is syncing again (for example, after a temporary connection loss or restart of the secondary) • ACTIVE: initialization or sync with primary is complete and secondary is continuously replicating. If crash occurs, no data loss occurs in SYNC replication mode. • ERROR: replication can't take place because the secondary system isn't accessible (details can be found in Replication Details)
Replication Details	Additional information for Replication Status, for example, the error text if status is ERROR.

Column	Description
Full Sync	Indicates if the service is currently operating in sync replication mode with the full sync option set. If full sync is enabled 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. In a second step, it's internally activated when the secondary is connected and becomes ACTIVE. • DISABLED: full sync isn't configured at all • ENABLED: full sync is configured, but it isn't yet active, so transactions don't block in this state. To become active, the secondary has to connect and Replication Status has to be ACTIVE. • ACTIVE: full sync mode is configured and active. If a connection of a connected secondary is getting closed, transactions on the primary side block in this state. If full sync is enabled when an active secondary is currently connected, the FULL_SYNC is immediately set to ACTIVE.
Secondary Fully Recoverable	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. FALSE: Log segments needed for a full database recovery are missing. After takeover, a full data backup has to be executed before a full recovery up to the most recent time point can be executed.
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 (only relevant for replication mode async; 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
Replayed Log Position	Log end position of the last known replayed log buffer on secondary site

Column	Description
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. • SYNC/SYNCMEM: total round-trip time to send the log buffers and receive the acknowledgment. • ASYNC: start time when sending the log buffers, and end time when the OS reports that the log buffers were sent (and the log shipping buffer space was freed). This time could be shorter than the SYNC/SYNCMEM duration
Time delay (ms)	Time delay between the last shipped log position time and the replayed log position time on the secondary
Size delay (Bytes)	Size delay between the last shipped log position size and the replayed log position size on the secondary (1 log position = 64 bytes)

Savepoints

Column	Description
Last Savepoint Version	Last savepoint version on primary
Last Savepoint Log Position	Log position of current savepoint
Last Savepoint Start Time	Timestamp of current savepoint
Last Shipped Savepoint Version	Last savepoint version shipped to secondary
Last Shipped Savepoint Log Position	Log position of last shipped savepoint
Last Shipped Savepoint Time	Timestamp of last shipped savepoint

Full Data Replica

Column	Description
Full Data Replica Shipped Count	Number of full data replicas shipped to secondary
Full Data Replica Shipped Total Size (Bytes)	Total size of all full data replica shipped to secondary
Full Data Replica Shipping Total Time (µs)	Duration for shipping all full data replica
Last Full Data Replica Shipped Size (Bytes)	Size of last full data replica shipped to secondary
Start Time of Last Full Data Replica	Start time of last full data replica
End Time of Last Full Data Replica	End time of last full data replica

Delta Data Replica

This information is only displayed if the operation mode is `delta_datashipping`.

Column	Description
Delta Data Replica Shipped Count	Number of delta data replicas shipped to secondary
Delta Data Replica Shipped Total Size (Bytes)	Total size of all delta data replicas shipped to secondary
Delta Data Replica Shipped Total Time (µs)	Duration for shipping of all delta data replicas
Size of Last Delta Data Replica (Bytes)	Size of last delta data replica
Start Time of Last Delta Data Replica	Start time of last data delta replica
End Time of Last Delta Data Replica	End time of last data delta replica

Log Shipping Backlog

Column	Description
Current Replication Backlog Size (Bytes)	Current replication backlog in bytes. The size of all log buffers that have been created on primary site, but not yet sent to the secondary site. Even in replication modes sync/syncmem this column can have a value different from 0. 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).
Max Replication Backlog Size (Bytes)	Max replication backlog in bytes (max value of BACKLOG_SIZE since system start).
Current Replication Backlog Time (µs)	Current replication backlog in microseconds. This time is the difference between time of the last sent log buffer and the current log buffer. 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 these buffers is the number of log buffers configured on primary site).
Max Replication Backlog Time (µs)	Max replication backlog in microseconds (max value of BACKLOG_TIME since system startup).

Log Replay

This information is only displayed if the operation mode is `logreplay` and `logreplay_readaccess`.

Column	Description
Replay Backlog Size (Bytes)	Specifies the size of all log buffers that have been shipped to the secondary site but haven't yet been replayed on the secondary site.
Max Replay Backlog Size (Bytes)	Specifies the maximum value of the REPLAY_BACKLOG_SIZE since the system startup.
Replay Backlog Time (µs)	Specifies the time difference between the time of the last shipped log buffer and the last replayed log buffer on the secondary site.
Max Replay Backlog Time (µs)	Specifies the maximum value of REPLAY_BACKLOG_TIME since the system startup.

Related Information

SQL and System View Reference

7.4.2 General Prerequisites for Configuring SAP HANA System Replication

Before you configure SAP HANA system replication, several prerequisites must be fulfilled.

- The primary and secondary systems are both installed and configured. You've verified that both are independently up and running.
- SAP HANA systems can only be replicated as the whole system, which 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 isn't possible.
- The configuration of active hosts in the primary and secondary system must be the same, which means that the number of active hosts, the names of the host roles, failover groups, and worker groups must be identical in both systems. So, if there's a standby host on the primary system it need not be available on the secondary system and vice versa.
- Check that the host names in the primary system are different to the host names used in the secondary system. You can see the SAP HANA host name for each host in the environment variable SAP_RETRIEVAL_PATH (/usr/sap/<SID>/HDB<InstNo>/<hostname>) and with the python script `landscapeHostConfiguration.py`. For more information, see *Host Name Resolution for System Replication* and *Checking the Status with landscapeHostConfiguration.py* in the *SAP HANA Administration Guide*.
  If the host names of the primary and the secondary system are the same (for example, because two systems are used that have identical host names), change the host names used on the secondary system. For more information, see *Rename an SAP HANA System Host* in the *SAP HANA Administration Guide*.
- System replication between two systems on the same host isn't supported.
- Both systems should run on the same endianness platform.
- You're logged on to both systems as the operating system user (user `<sid>`adm) or you've provided its credentials when prompted.
- You need the operating system user to set up a system replication landscape, to perform a takeover and a failback, as well as to disable system replication with the SAP HANA cockpit. For more information, see *Operating System User <sid>adm* and *Connect to a Database using Database Credentials*.
- The secondary system must have the same SAP system ID (<SID>) and `instance number` as the primary system.

  > **i Note**
  >
  > The primary system replicates all relevant license information to the secondary system. An additional license isn't required. For more information, see *SAP Note 2211663*.

- All configuration steps have to be executed on the master name server node only.
- The .ini file configuration must be similar for both systems. Any changes made manually, or by SQL commands on one system should be manually duplicated on the other system.
  Automatic configuration parameter checks alert you to configuration differences between the two systems.

> **i Note**
>
> To keep the ini file configuration similar on both systems, the INI parameter checker is per default configured to check for differences. Additionally, it can be configured to replicate parameter changes from the primary system to the secondary system.

- Ensure that log_mode is set to normal in the persistence section of the global.ini file. This mode means that the log segments are backed up.
- During an upgrade of the system replication landscape, the software version of the current secondary system must be equal or newer to the version of the current primary system.

  > **i Note**
  >
  > During a failback, the roles of your systems in the system replication landscape switch. Make sure in this case that your primary system doesn't have a newer software version than the secondary system.

  > **i Note**
  >
  > For Active/Active (read enabled) setups, the SAP HANA versions must be the same on the primary and the secondary system. Use this setup mainly during the upgrade process of the system replication landscape.

- To secure the system replication communication channel between the primary and the secondary system, configure the ini parameters `[system_replication_communication] / listeninterface` and `allowed_sender` as described in *Host Name Resolution for System Replication*.

  > **i Note**
  >
  > If you plan to add SAP HANA dynamic tiering to your landscape in the future, see*SAP Note 2447994* before you enable HANA system replication. SAP HANA dynamic tiering requires certain communication ports, operation modes, and replication modes.

- If a new tenant database is created in a running SAP HANA system replication, it must be backed up to participate in the replication. Afterward, 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 won't be operational after takeover and have to be recovered with backup and recovery. See the *SAP HANA Database Backup and Recovery* section of the *SAP HANA Administration Guide*.
- Before you configure SAP HANA system replication, you must copy the system PKI SSFS .key and the .dat file from the primary system to the secondary system:
  `/usr/sap/<SID>/SYS/global/security/rsecssfs/data/SSFS_<SID>.DAT`
  `/usr/sap/<SID>/SYS/global/security/rsecssfs/key/SSFS_<SID>.KEY`
  For more information, see *SAP Note 2369981 Required configuration steps for authentication with HANA System Replication*.
  If you installed XS advanced, you must also copy the XSA SSFS .key and the .dat file from the primary system to the secondary system in the following directories:
  `/usr/sap/<SID>/SYS/global/xsa/security/ssfs/data/SSFS_<SID>.DAT`
  `/usr/sap/<SID>/SYS/global/xsa/security/ssfs/key/SSFS_<SID>.KEY`
  For more information, see *SAP Note 2300936 Host Auto-Failover & System Replication Setup with SAP HANA extended application services, advanced model*.
  The copied files become active during system restart. Therefore, it's recommended to copy them when the secondary system is offline (for example, before registration).
- For SAP HANA system replication, a port offset value of 100 is configured to reserve ports for system replication communication. The port number of the replication port is calculated by adding the value for

this replication port offset to the internal port number of the corresponding service. Thus, although the same `<instance number>` is used for primary and secondary systems, the `<instance number>`+1 is reserved for both systems, because this port range is needed for system replication communication. For SAP HANA systems, this port offset is set to 10000 shifting the ports from the 3`<instance number>`00 to the 4`<instance number>`00 port range for the services. This is necessary in SAP HANA system replication with SAP HANA systems, because after 3`<instance number>`99 is reached new tenant databases allocate port numbers of the next higher instance number.

> **Note**
>
> To avoid interference with ephemeral ports it might be necessary to adjust the OS port range when using SAP HANA system replication in combination with SAP HANA tenant databases. On Linux this can be accomplished with the following command in the system startup script: `echo "50000 65535" > /proc/sys/net/ipv4/ip_local_port_range`.

- In preparation for maintenance tasks (for example, near zero downtime upgrades), configure a user in the local userstore under the SRTAKEOVER key. For more information, see *Configure a User Under the SRTAKEOVER Key* in the *SAP HANA Administration Guide*.
- SAP HANA dynamic tiering isn't supported with multitarget system replication. For more information about SAP HANA system replication with SAP HANA dynamic tiering, see *SAP Note 2447994*.

Related Information

Host Name Resolution for System Replication
Rename an SAP HANA System Host
Secure Internal Communication Between Sites in System Replication Scenarios
SAP HANA Database Backup and Recovery
Copying and Moving Tenant Databases
Configure a User Under the SRTAKEOVER Key
Connect to a Database With SSO or SAP HANA Credentials [page 97]
Operating System User <sid>adm
Checking the Status with landscapeHostConfiguration.py
SAP HANA System Replication
SAP Note 2211663
SAP Note 2369981
SAP Note 2300936
SAP Note 2447994

7.4.3 Configure SAP HANA System Replication

To configure SAP HANA system replication in the SAP HANA cockpit, first enable system replication on the primary system and then register the secondary system.

Prerequisites

Access the *Database Overview* by clicking *Database Directory* and then clicking on your database name.

- You've considered all the general prerequisites needed to set up system replication. For more information, see *General Prerequisites for Setting Up SAP HANA System Replication*.
- You've registered the system database for both systems in the SAP HANA cockpit.
- You have the operating system user to set up a system replication landscape with the SAP HANA cockpit. For more information, see *Operating System User <sid>adm* in the *SAP HANA Administration Guide* and *Connect to a Database using Database Credentials*.

Context

In the *Monitoring* or *All* view, on the *Database Overview* page, the *System Replication* card provides the possibility to configure system replication. Once the configuration is done, the card displays information on the operation mode, the replication mode, the configuration type, and the status of system replication.

The secondary system can be registered from the primary system or from the *Database Overview* page of the SAP HANA cockpit. You can register again a previously stopped secondary system when a full data shipping is needed or when you want to change the operation mode

Related Information

General Prerequisites for Configuring SAP HANA System Replication [page 173]
Connect to a Database With SSO or SAP HANA Credentials [page 97]
Operating System User <sid>adm
SAP HANA System Replication

7.4.3.1 Configure SAP HANA System Replication from the Primary System

To configure SAP HANA system replication, first enable system replication on the primary system and then register the secondary system. Use the SAP HANA cockpit to execute these separate steps in one configuration step from the primary system.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

- You've considered all the general prerequisites needed to set up system replication. For more information, see *General Prerequisites for Setting Up SAP HANA System Replication*.

Context

This topic describes how to configure system replication from the primary system in SAP HANA cockpit in one configuration step. You can use this method for 2-tier and 3-tier setups.

> **i Note**
>
> If you plan to add SAP HANA dynamic tiering to your landscape in the future, see *SAP Note 2447994* before you enable HANA system replication. SAP HANA dynamic tiering requires certain communication ports, operation modes, and replication modes.

Procedure

1. In the *Monitoring* or *All* view, on the *Database Overview* page of the system database (SYSTEMDB) of the future primary system, choose the *System Replication* card.

   If you never configured system replication before, this card displays the message *System replication is not yet enabled for this system*.

   The *System Replication* page opens. If you performed a data backup before enabling system replication, this page displays the last data backup on the top left and the *Configure System Replication* button on the top right.
2. Choose *Configure System Replication*.

   The *System Replication Configuration* dialog opens, allowing you to run the configuration in background.
3. Enter the logical name used to represent the primary system in the *Tier 1 System Details* screen area.
4. Enter the logical name used to represent the secondary system in the *Tier 2 System Details* screen area.

Keep in mind that the secondary system must have the same SAP system ID (<SID>) and instance number as the primary system so that they're identified as secondaries.

5. Select the secondary system host and mark the checkbox below this area to stop the system.
6. Select a replication mode. For more information on the available replication modes, see *Replication Modes for SAP HANA System Replication* in the *SAP HANA Administration Guide*.
7. Select an operation mode. For more information on the available operation modes, see *Operation Modes for SAP HANA System Replication* in the *SAP HANA Administration Guide*.
8. Decide whether to initiate a full data shipping or not.
9. Check *Start Secondary after Registration*.
10. Optional: To add a third tier to your system replication landscape configuration, click *Add Tier 3 System* on the bottom left.
11. Choose *Configure*.

Related Information

General Prerequisites for Configuring SAP HANA System Replication [page 173]
Replication Modes for SAP HANA System Replication
Operation Modes for SAP HANA System Replication
SAP HANA System Replication
SAP Note 2447994

7.4.3.2 Configure SAP HANA System Replication from the Primary and the Secondary Systems

To set up SAP HANA system replication, first enable system replication on the primary system and then register the secondary system. Use the SAP HANA cockpit to execute these configuration steps on the primary system and separately on the secondary system.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

- You've considered all the general prerequisites needed to set up system replication. For more information, see *General Prerequisites for Setting Up SAP HANA System Replication*.

Context

This topic describes how to enable system replication on the primary system and then register the secondary system using the SAP HANA cockpit. You can use this method to configure any system replication setups you want.

> **i** Note
>
> If you plan to add SAP HANA dynamic tiering to your landscape in the future, see *SAP Note 2447994* before you enable HANA system replication. SAP HANA dynamic tiering requires certain communication ports, operation modes, and replication modes.

Procedure

1. In the *Monitoring* or *All* view, on the *Database Overview* page of the system database (SYSTEMDB) of the future primary system, choose the *System Replication* card.

   If you never configured system replication before, this card displays the message *System replication is not yet enabled for this system*.

   The *System Replication* page opens. If you performed a data backup before enabling system replication, this page displays overview information on the primary system on the top left and the *Enable This System as Primary* link on the top right.
2. Enter the logical name used to represent the primary system and choose *Configure* on the bottom right.
3. On the *Database Overview* page of the future secondary system, click on the *Services* card.
4. Choose *Stop System* on the bottom right, because the system has to be offline in order to be registered as a secondary system.

   Back on the *Database Overview* page of the future secondary system, the *Services* card displays the status *Stopped*.
5. On the *Database Overview* page of the secondary system, choose the *System Replication* card.

   The *System Replication* page opens, displaying overview information about the secondary system on the top left and the *Register Secondary System* button on the top right.
6. Choose *Register Secondary System*.

   The *System Replication Configuration* page opens.
7. On the *System Replication Configuration* page, enter the logical name used to represent the secondary system.
8. On the *System Replication Configuration* page, select a replication mode. For more information on the available replication modes, see *Replication Modes for SAP HANA System Replication* in the *SAP HANA Administration Guide*.
9. Select an operation mode. For more information on the available operation modes, see *Operation Modes for SAP HANA System Replication* in the *SAP HANA Administration Guide*.
10. Enter the host of the source system.

> **Note**
>
> If you're operating a distributed system on multiple hosts, enter the name of the host on which the master name server is running.

11. Check *Start Secondary after Registration*.
12. Review the configured information and choose *Configure* on the bottom right.

    The *System Replication Configuration* dialog opens. After the configuration is complete, the *System Replication Overview* page displays information on the configured systems.

Related Information

General Prerequisites for Configuring SAP HANA System Replication [page 173]
Replication Modes for SAP HANA System Replication
Operation Modes for SAP HANA System Replication
SAP HANA System Replication
SAP Note 2447994

7.4.3.3 Reinitialize the Secondary System

You can register again a previously stopped secondary system using the SAP HANA cockpit.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

You can register again a previously stopped secondary system. Reinitialize when a full data shipping is needed or when you want to change the operation mode.

> **Note**
>
> The *System Replication* card isn't available in the tenant database.

Procedure

1. In the *Monitoring* or *All* view, on the *Database Overview* page of the system database (SYSTEMDB) of the stopped secondary system, choose the *System Replication* card.
2. On the *System Replication Overview*, choose *Reinitialize Secondary System* on the top right.
3. On the *System Replication Configuration* page, you can now change the configuration.
   Change the operation mode or resync the persistencies using the *Initiate full data shipping* option.

   The secondary system is up and running again.

Related Information

SAP HANA System Replication

7.4.3.4 Use Secondary Time Travel

You can start the secondary system or the log replay at a previous point in time.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

You have the `logreplay_readaccess` operation mode.

Context

Secondary time travel allows you to quickly access data that was deleted in the original system. You can start a secondary time travel on system database (SYSTEMDB) or SINGLEDB secondaries.

To prepare the system for time travel, the `global.ini/[system_replication]/timetravel_max_retention_time` parameter must be configured on the secondary system. This parameter defines the time period to which the secondary system can be brought back from the past. Optionally, the `global.ini/[system_replication]/timetravel_snapshot_creation_interval` parameter can be adapted to adjust the secondary's snapshot creation. After setting these parameters, the secondary starts retaining logs and keeping created snapshots.

The available snapshots for secondary time travel are visible on the *System Replication Overview*. Snapshots are kept on the secondary system for a defined time travel period to prepare the secondary system for time travel. These snapshots can be used later to start the system at an older point in time. Additional log will be

retained on the secondary system starting from the oldest time travel snapshot. After opening the old snapshot, the additional log has to be replayed to reach the requested point in time.

Procedure

1. In the *Monitoring* or *All* view, on the *Database Overview* page of your secondary system, select the *Manage system configuration* link to define the `global.ini/[system_replication]/timetravel_max_retention_time` and the `global.ini/[system_replication]/timetravel_snapshot_creation_interval` parameters.
2. Open the *System Replication* card of your secondary system.
3. On the *System Replication Overview* of your secondary system, choose *Start Time Travel* .

   The *Start Time Travel*dialog opens.
4. Select the date and the time to restart the secondary system and choose *Start*.

7.4.4 Perform a Takeover

You can perform a takeover on your secondary system using the SAP HANA cockpit.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

- It's recommended to stop the primary system before starting a takeover.

  > **i** Note
  >
  > If you're performing a takeover as part of a planned downtime, first make sure that the primary system has been fully stopped before performing a takeover to the secondary system.
- The secondary system must be fully initialized.
- You need the operating system user to perform a takeover with the SAP HANA cockpit. For more information, see *Operating System User <sid>adm* in the *SAP HANA Administration Guide* and *Connect to a Database using Database Credentials*.
- The takeover command can be executed both when the secondary system is offline or online.

Procedure

1. In the *Monitoring* or *All* view, on the *Database Overview* page of the system database (SYSTEMDB) of the secondary system meant to perform the takeover, choose the *System Replication* card.

The *System Replication* card opens displaying the *System Replication Overview*.

2. Choose *Take Over* on the top right.
3. To start the takeover, click *Start Takeover* in the *Takeover* dialog.

   You can also start a takeover with handshake by choosing to fully synchronize the secondary system. For more information about the takeover with handshake, see *Takeover with Handshake* in the *SAP HANA Administration Guide*.
4. Stop the primary system from the *Services* card on the *Database Overview*.

Related Information

Connect to a Database With SSO or SAP HANA Credentials [page 97]
Operating System User <sid>adm
Takeover with Handshake
SAP HANA System Replication

7.4.5 Perform a Failback

To perform a failback in the SAP HANA cockpit, register the former primary system as secondary to the current primary.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

- You need the operating system user to perform a failback with the SAP HANA cockpit. For more information, see *Operating System User <sid>adm* in the *SAP HANA Administration Guide* and *Connect to a Database using Database Credentials*.
- The former primary system isn't running.
- The current primary system is running.

Context

In the SAP HANA cockpit, you can perform a failback either from the current primary system or from the former primary system (the future secondary). The configuration steps are the same as described in *Configure System Replication from the Primary System* or *Configure System Replication from the Primary and the Secondary System*.

This procedure describes how to register the former primary as a new secondary. Use the *System Replication* card on the system *Database Overview* page of the former stopped primary to register this system as a new secondary.

Procedure

1. Register the secondary system as follows:
   a. In the *Monitoring* or *All* view, on the *Database Overview* page of the system database (SYSTEMDB) of the former primary system, choose the *System Replication* card.

      The *System Replication* page opens.
   b. Choose *Register as Secondary* on the top right.

      The *Register Secondary System* page opens.
   c. On the *Register Secondary System* page, enter the logical name used to represent the secondary system.
   d. Select a replication mode. For more information on the available replication modes, see *Replication Modes for SAP HANA System Replication*.
   e. Select an operation mode. For more information on the available operation modes, see *Operation Modes for SAP HANA System Replication*.
   f. Enter the host of the source system.

      > **i Note**
      >
      > If you're operating a distributed system on multiple hosts, enter the name of the host on which the master name server is running.

   g. Check *Start Secondary after Registration*.
2. Review the configured information and choose *Configure* on the bottom right.

Results

The original primary system is now registered as the secondary system with the current primary system (that is, the original secondary system). The secondary system is getting in sync again with the primary system. As such, it's attempting to avoid a full data shipping.

Verify that the secondary system replication status is `All services are active and in sync`.

Related Information

Perform a Failback [page 183]
Configure SAP HANA System Replication from the Primary and the Secondary Systems [page 178]
Replication Modes for SAP HANA System Replication
Operation Modes for SAP HANA System Replication

Connect to a Database With SSO or SAP HANA Credentials [page 97]
Operating System User <sid>adm
SAP HANA System Replication

7.4.6 Disable SAP HANA System Replication

You can disable SAP HANA system replication in an SAP HANA system using the SAP HANA cockpit.

Prerequisites

- The secondary system must be offline.
- You need the operating system user to disable system replication with the SAP HANA cockpit. For more information, see *Operating System User <sid>adm* in the *SAP HANA Administration Guide* and *Connect to a Database using Database Credentials*.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. In the *Monitoring* or *All* view, on the *Database Overview* page of the system database (SYSTEMDB), use the *Services* card .
2. Unregister the secondary system as follows:
   a. In the *Monitoring* or *All* view, on the *Database Overview* page of the system database (SYSTEMDB) of the stopped secondary system, choose the *System Replication* card to stop the secondary system.

      The *System Replication* card opens displaying the *System Replication Overview*.
   b. Choose *Unregister this Secondary* on the top right.

      Depending whether your system should be online or offline after unregistering it, check the *Start system after unregistration* option in the confirmation dialog and choose *OK*. For more information, see *SAP Note 1945676*.
3. Disable system replication on the primary system as follows:
   a. In the *Monitoring* or *All* view, on the *Database Overview* page of the primary system, choose the *System Replication* card.

      The *System Replication* card opens displaying the *System Replication Overview*.
   b. Choose *Disable System Replication* on the top right and confirm that you want to disable system replication.

      The *Ignore the secondary system* option allows you to disable the primary system even though the secondary is still attached. This could be relevant, if the secondary has been uninstalled in the meantime.

Related Information

Connect to a Database With SSO or SAP HANA Credentials [page 97]
Operating System User <sid>adm
SAP HANA System Replication
SAP Note 1945676

7.5 Admission Control

You can use SAP HANA cockpit to manage peak load by applying processing limits and determining how to handle new requests if the system is close to the point of saturation.

You can apply thresholds to define an acceptable limit for the percentage of memory usage or percentage of CPU capacity. Threshold values for admission control to determine when requests are queued or rejected are defined as configuration parameters. See also *Managing Peak Load (Admission Control)* in the *SAP HANA Administration Guide*.

New requests are queued until adequate processing capacity is available or a timeout is reached. Also, a higher threshold can be defined to determine the maximum workload level above which new requests are rejected. If requests have been queued, items in the queue are processed when the load on the system reduces below the threshold levels. If the queue exceeds a specified size, or if items are queued for longer than a specified period of time, they're rejected.

Related Information

Manage Admission Control [page 186]
Managing Peak Load (Admission Control)

7.5.1 Manage Admission Control

Manage your peak load by determining whether new, incoming statement requests are rejected or queued based on memory and CPU statistics.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Admission control checks the current resource consumption within a system and, along with the defined threshold values, decides whether or not a new statement is admitted during peak situations in the system.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click on the *Admission Control* card.
2. Click *Manage Admission Control* to configure CPU and memory thresholds.
3. Ensure you have checked *Enable admission control so that requests can be rejected or queued based on CPU and memory resource tracking statistics.*
4. View or adjust the default settings for each of the parameters and thresholds.
5. Select *Save*.
6. (Optional) Click *Monitor* to switch over to the *Admission Control Monitor* to view advanced admissions control statistics.

Results

Any changes you make are saved as new settings in `global.ini [session_admission_control]`, or, for admission control log management settings, in `global.ini [session_admission_control_events]`. You can view these new settings using the *Configuration Manager*.

Related Information

Database Configuration in SAP HANA Cockpit [page 443]
Admission Control Default Configuration Parameters [page 188]

7.5.2 Manage Admission Control Advanced Analytics

Use the *Admission Control Monitor* to view advanced admissions control statistics.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

The *Admission Control Monitor* allows you to monitor the following:

- Queued, rejected, and admitted requests over time
- Event details for submitted requests
- Queued session requests

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click the *Admission Control* card.

   The *Admission Control Monitor* opens.
2. *Statistics*On the tab, filter session request counts according to *Host:Port* and by time period.

   You can view the session request counts statistics in either chart or table form and customize both according to your preferences using the *Settings* icon.

   Clicking anywhere on the *Admitted* line graph shows you the number of admitted sessions for the selected time.
3. On the *Events* tab, filter and customize the *Admissions Control Events* table to show only specific types of events, or only events from a specific *Host:Port* or *Connection ID*, over a specified amount of time.
   a. Click the *More* button next to a SQL statement to view the full SQL statement and click *Copy* to copy it to your clipboard.
4. On the *Active Queue* tab, view queued session requests by *Host:Port* or *Connection ID*.

Related Information

Managing Peak Load (Admission Control)
M_ADMISSION_CONTROL_STATISTICS System View
M_ADMISSION_CONTROL_EVENTS System View
M_ADMISSION_CONTROL_QUEUES System View

7.5.3 Admission Control Default Configuration Parameters

The default configuration parameters for admission control can be modified through SAP HANA Cockpit using *Admission Control Monitor*.

Changes are saved to `global.ini [session_admission_control]`, or, in the case of admission control log management settings, to `global.ini [session_admission_control_events]`.

Field Name	Parameter	Default	Min/Max	Detail
Enable admission control ...	enable	True		Enables or disables admission control.
CPU threshold for queuing new requests	queue_cpu_threshold	90	0/100	The percentage of CPU usage above which requests are queued.
Memory threshold for queuing new requests	queue_memory_threshold	90	0/100	The percentage of memory usage above which requests are queued.
CPU threshold for rejecting new requests	reject_cpu_threshold	0	0/100	The percentage of CPU usage above which requests are rejected. The default value 0 means that no requests are rejected, but can be queued.
Memory threshold for rejecting new requests	reject_memory_threshold	0	0/100	The percentage of memory usage above which requests are rejected. The default value 0 means that no requests are rejected, but can be queued.

Admission control log management

Field Name	Parameter	Default	Min/Max	Detail
Event logging level	event_level	2	1/5	Level of report event, where 1=OFF 2=BASIC (REJECT, QUEUE) 3=BASIC_CONFIG (REJECT, QUEUE, CONFIGURATION CHANGE) 4=DEBUG (REJECT, QUEUE, CONFIGURATION CHANGE, EXCLUDE) and 5=ALL (REJECT, QUEUE, CONFIGURATION CHANGE, EXCLUDE, ADMIT)
Add to log if wait time in queue is greater than	queue_wait_time_threshold	100000	0/ uint64_max	The length of time measured in microseconds for which a request must be queued above which it's included in the event log (default is one tenth of a second). If the parameter is set to 0, then events aren't logged.
Maximum number of log records	record_limit	1000000	1/ int32_max	The maximum record count permitted in the monitor of historical events.

Queue management

Field Name	Parameter	Default	Min/Max	Detail
Maximum queue size	max_queue_size	10000	1/int32_max	The maximum number of requests that can be queued. Requests above this number are rejected.
Interval for checking to release queued requests	dequeue_interval	1000	100/int32 max	Unit: milliseconds. Use this parameter to set the frequency of the check to re-evaluate the load in comparison to the thresholds. The default is 1000 ms (1 second). This value is recommended to avoid overloading the system, though values from 100 ms are supported.
Batch size when releasing queued requests	dequeue_size	50	1/9999	Use this parameter to set the de-queue batch size, which is the number of queued items that are released together once the load is sufficiently reduced. This value can be from 1 through 100 queued requests.

Statistics collection management

Field Name	Parameter	Default	Min/Max	Detail
Averaging factor	averaging_factor	70	1/100	This percentage value gives a weighting to the statistic averaging process. A low value has a strong moderating effect (but may not adequately reflect real CPU usage). A value of 100% means that no averaging is performed, so only the current value for memory and CPU consumption is considered.
Statistics collection interval	statistics_collection_interval	1000	100/int32 max	Unit milliseconds. The statistics collection interval is set by default to 1000 ms (1 second) which has a negligible effect on performance. Values from 100 ms are supported. Statistics details are visible in the view M_ADMISSION_CONTROL_STATISTICS.

7.6 Monitoring Multi-Host Systems

Monitor the system health and network traffic between hosts.

7.6.1 Monitoring the Network Between Multiple Hosts

For scale-out systems, it's possible to monitor network traffic between hosts using the *Monitor Network* link in the SAP HANA cockpit.

In the *Monitoring* or *All* view, on the *Database Overview* page, click *Network Monitor* on the *Monitoring* card. On the *Monitor Network* page, you can view the number of hosts and use the following tabs to monitor the network for multiple hosts:

- *Network Traffic*
  Use this tab to understand the role of each host and the size of the sent (*Request Size*) and received data (*Response Size*) between the hosts of the scale-out SAP HANA database. The sender host sends requests to the receiver host which responses. You can change the unit on the top right.
- *Network Speed Check (Internal Communication)*
  The list offers an overview of all network channels between the involved hosts starting with the slowest network connection.
  The *Measure Network Speed* link offers the possibility to measure the network speed between the hosts in a scale-out SAP HANA database. You can select the size of the package for the speed check.
- *Network Speed Check (System Replication Communication)*
  The list offers an overview of all network channels between the involved hosts in the system replication configuration.
  The *Measure Network Speed* link offers the possibility to measure the network speed between the hosts in a system replication configuration.

7.6.2 Monitoring Multi-Host Health

For scale-out or multiple-host systems, SAP HANA cockpit provides status information on the health of database components on their respective servers, and on resource utilization of hardware components, including CPU, memory, network, and storage on the respective servers.

On a multi-host database, in the *Monitoring* or *All* view, on the *Database Overview* page, click *Monitor Multi-Host Health* on the *Monitoring* card.

The multi-host health charts display the most recent 10 minutes of data, organized by hosts within the system. The host types include master, standby, dynamic tiering, streaming. The health metrics are:

Metric	Details
Status	Status of the main service (indexserver)
Critical alerts	Number of high priority alerts for the given host

Metric	Details
CPU %	Percentage of CPU usage for the given host
Memory %	Percentage of memory usage for the given host
Out of Memory Events	Number of out of memory events
Unloads	Number of columns unloaded due to low memory
Disk Usage %	Percentage of disk usage for the given host
Network I/O	Number of network input/output events
Statements	Changes in number of statements per second
Versions	Number of multi-version concurrency control versions
Disk I/O	Number of disk input/output events

To reduce the number of hosts or show particular hosts side by side, click the ⚙ *(Settings)* icon and select the desired hosts from the list.

By selecting specific multi-host health information, you can drill down to details about alerts when critical alerts are present.

When you click on a chart, instead of navigating directly to the Performance Monitor, Memory Analysis, or Workload Analysis page, a dialog box appears that re-displays the chart, and offers multiple actions.

Clicking on...	Gives the option to...
CPU chart	• *View CPU chart* displays the Performance Monitor with a CPU chart • *View all metrics for this host* displays the Performance Monitor with all metrics • *View CPU for all hosts* displays the Performance Monitor with a CPU chart for each host in the scale-out database
Memory chart	• *Analyze memory history* displays Memory Analysis • *View memory chart* displays the Performance Monitor with a memory chart • *View all metrics for this host* displays the Performance Monitor with all metrics • *View memory for all hosts* displays the Performance Monitor with a memory chart for each host in the scale-out database.
Column Unload chart	• *View column unload chart* displays the Performance Monitor with a column unload chart • *View all metrics for this host*displays the Performance Monitor with all metrics • *View column unloads for all hosts*displays the Performance Monitor with a column unload chart for each host in the scale-out system

Clicking on...	Gives the option to...
Disk I/O chart	• *View column unload chart* displays the Performance Monitor with a single chart containing: Data Write Size, Data Read Size, Log Write Size, Log Read Size, Data Backup Write Size, Data Backup Read Size • *View all metrics for this host* displays the Performance Monitor with all metrics • *View disk I/O for all hosts*displays the Performance Monitor with a chart (containing the same metrics as the preceding single chart) for each host in the scale-out system
Disk Usage chart	• *View disk usage chart* displays the Performance Monitor with a disk used chart • *View all metrics for this host* displays the Performance Monitor with all metrics • *View disk usage for all hosts* displays the Performance Monitor with a disk used chart for each host in the scale-out system
MVCC Versions chart	• *View MVCC versions chart*displays the Performance Monitor with a single chart containing: Acquired Record Locks, Active Commit ID Range, Active Versions • *View all metrics for this host* displays the Performance Monitor with a shows all metrics • *View MVCC versions for all hosts*displays the Performance Monitor with a chart (containing the same metrics as the preceding single chart) for each host in the scale-out system
Statements chart	• *Analyze workload* displays Workload Analysis • *View statements chart* displays the Performance Monitor with a statements chart • *View all metrics for this host* displays the Performance Monitor with all metrics • *View statements for all hosts*displays the Performance Monitor with a statements chart for each host in the scale-out system
Network I/O chart	• *View network I/O chart* displays the Performance Monitor with a single chart containing Network In, Network Out • *View all metrics for this host* displays the Performance Monitor with all metrics • *View network I/O for all hosts* displays the Performance Monitor with a chart containing Network In, Network Out for each host in the scale-out system

Related Information

The Performance Monitor [page 515]

7.7 Disk Usage: Monitor Disk Volume

In order to ensure that the database can always be restored to its most recent committed state, you can use the SAP HANA cockpit to check disk statistics to check that there's enough space on disk for data volumes and log volumes.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

A disk has multiple volumes. Each volume has a data volume and a log volume. Data volumes have one file (datavolume_0000.dat). Log volumes often have hundreds of files (multiple logsegment_000_0000003.dat; single logsegment_000_directory.dat). Log segment files have a state (Formatting, Preallocated, Writing, Closed, Truncated, BackedUp, RetainedFree, Free). Only log segment files with state Free can be reused. Log segment files have a fixed size although the size can vary per service. (For example, indexserver=1024MB; xsengine=8MB).

You may wish to monitor volume if, for example:

- You receive an alert about disk I/O read failure and want to see which volume has the issue and why.
- You're running out of disk space.
- You know that you're having a backup issue or a replication issue and want to understand how it's affecting disk usage.

Procedure

1. In the *Monitoring* or *All* view, on the *Database Overview* page, click the ... *(More)* icon on the *Disk Usage* card.
2. Select *Disk Volume Monitor*.
3. (Optional) Filter the information displayed in the chart and the table:
    - Using the arrow beside the title (top left), select a predefined variant, or manage and save a custom variant.
    - Select from the dropdown lists to display a specific combination of host, tenant (if applicable), volume type (data volume or log volume), service, and volume ID.
4. View the information in the table:

Column	Description
Volume ID	The ID of the volume.
Service	The name of the service.
Type	Whether this volume is a data volume or log volume.
Size [MB]	Current size of the volume.
Used [MB]	Amount of disk space used on the host's hard disk.
Used [%]	Amount of disk space used on the host's hard disk as a percentage of the whole.
State	For log files, whether the state is Formatting, Preallocated, Writing, Closed, Truncated, BackedUp, RetainedFree or Free. The log segment file's state indicates its availability for reuse.
Files	Number of files of the same type and state for the particular host and service.
Path	Location of the service's data and log files in the file system.
Host	The name of the host.
Tenant	The name of the tenant.

5. Drill down on a row to see *Disk Volume Details* for a specific volume.
6. Select the various tabs to move to the corresponding section of the volume details.

Section	Description
Data Volume Files	Displays the data volume file names as well as the size of each file and how much of it's currently in use, both in MB and as a percentage of its total size. Used size is the amount of data in the file. As the size of the file is automatically increased with the payload but not automatically decreased, used size and total size can be different.
Log Files	Displays log file names, total size (which, for log files, is equivalent to used size), and state. When a file is full, log entries are written to the next log segment file available.
Volume I/O Statistics	Displays aggregated I/O statistics for the volume, and, for comparison, other volumes in the system, in your choice of time periods: ○ Since the service was restarted (default) ○ Since the last manual reset You can reset the statistics collection for all volumes by selecting *Reset Volume I/O Total Statistics*.
Data Volume Page Statistics	Displays statistics on the data volume's pages (or blocks) broken down according to page size class. Superblocks are partitions of the data volume that contain pages of the same page size class. You can analyze how many superblocks are used for the specific size class and also how many pages/blocks are used. The fill ratio enables you to decide whether or not it makes sense to reorganize and release unnecessary superblocks, in other words, shrink the data volume.

7. In the *Volume I/O Statistics* section, use the left-most dropdown menu to display specific statistics associated with the volume:
   - Volume Size & Time
   - Volume Configuration
   - Advanced Write Statistics
   - Advanced Read Statistics

Related Information

Reclaim Space [page 196]
Persistent Data Storage in the SAP HANA Database
Monitoring Disk Space

7.7.1 Reclaim Space

You can reclaim space by reclaiming freed log segments and unused space in data volumes.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Perform a backup before reclaiming space.

Context

The reclaim operation is across hosts & volumes per database & tenant. You can reclaim all log files in Free state but at least `log_preformat_segment_count` segments (by default, two) per database service aren't reclaimed. For data volumes, the reclaim is calculated by volume size – (used size * specified `percentage_of_overload_size`).

Procedure

1. In the *Monitoring* or *All* view, on the *Database Overview* page, click the ... *(More)* icon on the *Disk Usage* card.
2. Select *Disk Volume Monitor*.
3. Select *Reclaim Space* from the top toolbar.
4. In the dialog, select *Reclaim (Free) log segments* and/or *Reclaim data volume*.
5. Select the *Reclaim Space* button.

8 Security and User Management View

This view helps you monitor and administer security and user management tasks on your database.

8.1 Security

Monitor critical security settings and perform administration tasks related to user authentication, data and communication encryption, and audit logging.

> **i Note**
>
> In addition to this documentation, please also refer to the *SAP HANA Security Guide* and *SAP HANA Security Checklists and Recommendations*.

> **i Note**
>
> Not all security-related administration and configuration tasks in SAP HANA are possible with SAP HANA cockpit. For some tasks, you may require other tools. For complete information about all administration tasks and tools, see the *SAP HANA Security Administration* section of the *SAP HANA Administration Guide*.

8.1.1 SAP HANA Security Checklists and Recommendations

SAP HANA has many configuration settings that allow you to customize your system. Some of these settings are important for the security of your system, and misconfiguration could leave your system vulnerable.

The checklists offer recommendations and information about optimizing your security configuration to help you run your SAP HANA securely. However, please note the following:

- The checklists and recommendations offered here are not exhaustive.
  In addition, depending on your specific implementation scenario and technical environment, some of the recommendations may not apply or be different.
- Do not use the checks as instructions on how to configure individual settings.
  If a particular check result indicates an insecure setting, refer to the indicated documentation and follow the instructions there to change the configuration setting.
- This document does not replace the *SAP HANA Security Guide*, the central document for all information relating to the secure operation and configuration of SAP HANA.

To view the security checklist from SAP HANA cockpit, on the *Database Overview* page, with the *Security and User Management* or *All* view selected, choose the *Security Checklist* link on the *Security Related Links* card.

Related Information

SAP HANA Security Guide

8.1.1.1 General Recommendations

General recommendations for keeping SAP HANA secure.

- Create a security concept for the SAP HANA scenario that you want to implement as early as possible in your implementation project.
- Install SAP HANA revisions that are marked as security-relevant as soon as possible. Do this by checking SAP HANA security notes either directly, or using services provided by SAP Support.
  For more information, see *SAP HANA Security Patches* in the SAP HANA Security Guide.

Related Information

SAP HANA Security Patches

8.1.1.2 Checklist for Secure Handover

If you received your SAP HANA system pre-installed from a hardware or hosting partner, there are several things we strongly recommend you do immediately after handover.

- Change the password of all operating system users, in particular the following:
  - `<sid>adm`
  - `root`
  - `sapadm`

  For more information, see your operating system documentation.
- In all databases, review all database users created by the installing party, and delete or deactivate those that are not needed in your scenario.

  > → Remember
  >
  > If you received a system with tenant databases, make sure to do this in all tenant databases and in the system database.

  For more information about database users that are created in the SAP HANA database by default, see the *SAP HANA Security Guide*.
- In all databases, change the password of all predefined database users, in particular the password of the database user `SYSTEM`. In addition, deactivate the `SYSTEM` user. For more information, see the *SAP HANA Security Guide*.

> → Remember
>
> If you received a system with tenant databases, make sure to do this in all tenant databases and in the system database.

> i Note
>
> Predefined internal technical users (`SYS`, `_SYS_*` users) are permanently deactivated and cannot be used to log on. It is not possible to change the password of these users.

- Change the following encryption master keys:
  - Instance secure store in the file system (SSFS)
  - System public key infrastructure (PKI) SSFS

  For more information about how to change the encryption master keys, see *SAP Note 2183624 (Potential information leakage using default SSFS master key in HANA)* and the *SAP HANA Administration Guide*.
- Re-create the system public key infrastructure (PKI) used to protect internal communication in order to create new certificates and private keys. You can trigger this by deleting the system PKI SSFS. Alternatively, you can use SAPControl to reset the system PKI with the methods `UpdateSystemPKI[<force>]` and `UpdateInstancePSE[<force>]`.

> i Note
>
> In a system replication landscape, you must copy the system PKI SSFS data file and key file from the primary system to the same location on the secondary system(s). For more information, see the section on secure internal communication in the *SAP HANA Security Guide*.

Related Information

Deactivate the SYSTEM User
Change a Database User [page 361]
Change the SSFS Master Keys [page 233]
SAP Control WebService
Secure Internal Communication Between Sites in System Replication Scenarios
SAP Note 2183624

8.1.2 SAP HANA Database Checklists and Recommendations

Checklists and recommendations to help you operate and configure the SAP HANA database securely

> → Tip
>
> SAP Note 1969700 contains collections of useful SQL statements for monitoring and analyzing the SAP HANA database. The statements contained in the file `HANA_Security_MiniChecks.txt` perform all of the SQL-based checks listed in this documentation.

8.1.2.1 Recommendations for Database Users, Roles, and Privileges

Recommendations for securing access to SAP HANA.

SYSTEM User

Default	The database user `SYSTEM` is the most powerful database user with irrevocable system privileges. The `SYSTEM` user is active after database creation.
Recommendation	Use `SYSTEM` to create database users with the minimum privilege set required for their duties (for example, user administration, system administration). Then deactivate `SYSTEM`. You may however temporarily reactivate the `SYSTEM` user for emergency or bootstrapping tasks. See *Deactivate the SYSTEM User* in the *SAP HANA Security Guide*. **i Note** The `SYSTEM` user is not required to update the SAP HANA database system; a lesser-privileged user can be created for this purpose. However, to upgrade SAP support package stacks, SAP enhancement packages and SAP systems using the Software Update Manager (SUM) and to install, migrate, and provision SAP systems using the Software Provisioning Manager (SWPM), the `SYSTEM` user **is required** and needs to be temporarily reactivated for the duration of the upgrade, installation, migration or provisioning.
How to Verify	In the system view `USERS`, check the values in columns `USER_DEACTIVATED`, `DEACTIVATION_TIME`, and `LAST_SUCCESSFUL_CONNECT` for the user `SYSTEM`.
Related Alert	No
More Information	See the sections on predefined users and deactivating the SYSTEM user in the *SAP HANA Security Guide*.

Password Lifetime of Database Users

Default	With the exception of internal technical users (`_SYS_*` users), the default password policy limits the lifetime of user passwords to 182 days (6 months).

Recommendation	Do not disable the password lifetime check for database users that correspond to real people. In 3-tier scenarios with an application server, only technical user accounts for the database connection of the application server should have a password with an unlimited lifetime (for example, SAP<sid> or DBACOCKPIT). **i Note** Such technical users should have a clearly identified purpose and the minimum authorization required in SAP HANA.
How to Verify	In the USERS system view, check the value in the column IS_PASSWORD_LIFETIME_CHECK_ENABLED. If it is FALSE, the password lifetime check is disabled. The time of the last password change is indicated in the column LAST_PASSWORD_CHANGE_TIME.
Related Alert	No
More Information	See the section on the password policy in the *SAP HANA Security Guide*.

System Privileges

Default	System privileges authorize database-wide administration commands. The users SYSTEM and _SYS_REPO have all these privileges by default.

Recommendation	System privileges should only ever be granted to users that actually need them. In addition, several system privileges grant powerful permissions, for example, the ability to delete data and to view data unfiltered and should be granted with extra care as follows: Only administrative or support users should have the following system privileges in a production database: • `CATALOG READ` • `TRACE ADMIN` In a database of any usage type, the following system privileges should be granted only to administrative users who actually need them: • `ADAPTER ADMIN` • `AGENT ADMIN` • `AUDIT ADMIN` • `AUDIT OPERATOR` • `BACKUP ADMIN` • `BACKUP OPERATOR` • `CERTIFICATE ADMIN` • `CREATE REMOTE SOURCE` • `CREDENTIAL ADMIN` • `ENCRYPTION ROOT KEY ADMIN` • `EXTENDED STORAGE ADMIN` • `INIFILE ADMIN` • `LDAP ADMIN` • `LICENSE ADMIN` • `LOG ADMIN` • `MONITOR ADMIN` • `OPTIMIZER ADMIN` • `RESOURCE ADMIN` • `SAVEPOINT ADMIN` • `SERVICE ADMIN` • `SESSION ADMIN` • `SSL ADMIN` • `TABLE ADMIN` • `TRUST ADMIN` • `VERSION ADMIN` • `WORKLOAD ADMIN` • `WORKLOAD * ADMIN`
How to Verify	To check which user has a particular system privilege, query the `EFFECTIVE_PRIVILEGE_GRANTEES` system view, for example: `SELECT * FROM EFFECTIVE_PRIVILEGE_GRANTEES WHERE OBJECT_TYPE = 'SYSTEMPRIVILEGE' AND PRIVILEGE = 'SSL ADMIN' AND GRANTEE NOT IN ('SYSTEM','_SYS_REPO');`
Related Alert	No
More Information	See the section on system privileges in the *SAP HANA Security Guide* and the section on system views for verifying user authorization in the *SAP HANA Administration Guide*.

System Privileges: Critical Combinations

Default	The users `SYSTEM` and `_SYS_REPO` have all system privileges by default.
Recommendation	Critical combinations of system privileges should not be granted together, for example: • `USER ADMIN` and `ROLE ADMIN` • `CREATE SCENARIO` and `SCENARIO ADMIN` • `AUDIT ADMIN` and `AUDIT OPERATOR` • `CREATE STRUCTURED PRIVILEGE` and `STRUCTUREDPRIVILEGE ADMIN`
How to Verify	To check a user's privileges, query the `EFFECTIVE_PRIVILEGES` system view, for example: `SELECT * FROM "PUBLIC"."EFFECTIVE_PRIVILEGES" WHERE USER_NAME = '<USER_NAME>';`
Related Alert	No
More Information	See the section on system privileges in the *SAP HANA Security Guide* and the section on system views for verifying user authorization in the *SAP HANA Administration Guide*.

System Privilege: DATA ADMIN

Default	The system privilege `DATA ADMIN` is a powerful privilege. It authorizes a user to execute all data definition language (DDL) commands in the SAP HANA database. Only the users `SYSTEM` and `_SYS_REPO` have this privilege by default.
Recommendation	No user or role in a production database should have this privilege.
How to Verify	You can verify whether a user or role has the `DATA ADMIN` privilege by executing the statement: `SELECT * FROM EFFECTIVE_PRIVILEGE_GRANTEES WHERE OBJECT_TYPE = 'SYSTEMPRIVILEGE' AND PRIVILEGE = 'DATA ADMIN' AND GRANTEE NOT IN ('SYSTEM','_SYS_REPO');`
Related Alert	No
More Information	See the section on system privileges in the *SAP HANA Security Guide* and the section on system views for verifying user authorization in the *SAP HANA Administration Guide*.

System Privilege: DEVELOPMENT

Default	The system privilege `DEVELOPMENT` authorizes some internal `ALTER SYSTEM` commands. By default, only the users `SYSTEM` and `_SYS_REPO` have this privilege.
Recommendation	No user or role in a production database should have this privilege.

How to Verify	You can verify whether a user or role has the `DEVELOPMENT` privilege by executing the statement: `SELECT * FROM EFFECTIVE_PRIVILEGE_GRANTEES WHERE OBJECT_TYPE = 'SYSTEMPRIVILEGE' AND PRIVILEGE = 'DEVELOPMENT' AND GRANTEE NOT IN ('SYSTEM','_SYS_REPO');`
Related Alert	No
More Information	If requested by SAP HANA support, this privilege can be granted using SQL. It is **not** included in the privilege handling overview in the SAP HANA Security Guide. See the section *System Views for Verifying Users' Authorization* in the *SAP HANA Administration Guide*.

Analytic Privilege: _SYS_BI_CP_ALL

Default	The predefined analytic privilege `_SYS_BI_CP_ALL` potentially allows a user to access all the data in activated views that are protected by XML-based analytic privileges, regardless of any other XML-based analytic privileges that apply. Only the predefined roles `CONTENT ADMIN` and `MODELING` have the analytic privilege `_SYS_BI_CP_ALL` by default. By default, only the user `SYSTEM` has these roles.
Recommendation	Do not grant this privilege to any user or role in a production database.
How to Verify	You can verify whether a user or role has the `_SYS_BI_CP_ALL` privilege by executing the statement: `SELECT * FROM EFFECTIVE_PRIVILEGE_GRANTEES WHERE OBJECT_TYPE = 'ANALYTICALPRIVILEGE' AND OBJECT_NAME = '_SYS_BI_CP_ALL' AND PRIVILEGE = 'EXECUTE' AND GRANTEE NOT IN ('SYSTEM','MODELING', 'CONTENT_ADMIN');`
Related Alert	No
More Information	See the sections on privileges and predefined database roles in the *SAP HANA Security Guide* and the section on system views for verifying user authorization in the *SAP HANA Administration Guide*.

Debug Privileges

Default	No user has debug privileges
Recommendation	The privileges `DEBUG` and `ATTACH DEBUGGER` should not be assigned to any user for any object in production systems.

How to Verify	You can verify whether a user or role has debug privileges by executing the statements: `SELECT * FROM GRANTED_PRIVILEGES WHERE PRIVILEGE='DEBUG' OR PRIVILEGE='ATTACH DEBUGGER';`
Related Alert	No
More Information	See the section on privileges in the *SAP HANA Security Guide* and the section on system views for verifying user authorization in the *SAP HANA Administration Guide*.

Predefined Catalog Role CONTENT_ADMIN

Default	The role `CONTENT_ADMIN` contains all privileges required for working with information models in the repository of the SAP HANA database. The user `SYSTEM` has the role `CONTENT_ADMIN` by default.
Recommendation	Only the database user used to perform system updates should have the role `CONTENT_ADMIN`. Otherwise do not grant this role to users, particularly in production databases. It should be used as a role template only.
How to Verify	You can verify whether a user or role has the `CONTENT_ADMIN` role by executing the statement: `SELECT * FROM GRANTED_ROLES WHERE ROLE_NAME = 'CONTENT_ADMIN' AND GRANTEE NOT IN ('SYSTEM');`
Related Alert	No
More Information	See the section on predefined database roles in the *SAP HANA Security Guide* and the section on system views for verifying user authorization in the *SAP HANA Administration Guide*.

Predefined Catalog Role MODELING

Default	The role `MODELING` contains the predefined analytic privilege `_SYS_BI_CP_ALL`, which potentially allows a user to access all the data in activated views that are protected by XML-based analytic privileges, regardless of any other XML-based analytic privileges that apply. The user `SYSTEM` has the role `MODELING` by default.
Recommendation	Do not grant this role to users, particularly in production databases. It should be used as a role template only.
How to Verify	You can verify whether a user or role has the `MODELING` role by executing the statement: `SELECT * FROM GRANTED_ROLES WHERE ROLE_NAME ='MODELING' AND GRANTEE NOT IN ('SYSTEM');`
Related Alert	No
More Information	See the section on predefined database roles in the *SAP HANA Security Guide* and the section on system views for verifying user authorization in the *SAP HANA Administration Guide*.

Predefined Catalog Role SAP_INTERNAL_HANA_SUPPORT

Default	The role `SAP_INTERNAL_HANA_SUPPORT` contains system privileges and object privileges that allow access to certain low-level internal system views needed by SAP HANA development support in support situations. No user has the role `SAP_INTERNAL_HANA_SUPPORT` by default.
Recommendation	This role should only be granted to SAP HANA development support users for their support activities.
How to Verify	You can verify whether a user or role has the `SAP_INTERNAL_HANA_SUPPORT` role by executing the statement: `SELECT * FROM EFFECTIVE_ROLE_GRANTEES WHERE ROLE_NAME = 'SAP_INTERNAL_HANA_SUPPORT';`
Related Alert	ID 63 (Granting of SAP_INTERNAL_HANA_SUPPORT role)
More Information	See the section on predefined database roles in the *SAP HANA Security Guide* and the section on system views for verifying user authorization in the *SAP HANA Administration Guide*.

Predefined Repository Roles

Default	SAP HANA is delivered with a set of preinstalled software components implemented as SAP HANA Web applications, libraries, and configuration data. The privileges required to use these components are contained within repository roles delivered with the component itself. The standard user `_SYS_REPO` automatically has all of these roles. Some may also be granted automatically to the standard user `SYSTEM` to enable tools such as the SAP HANA cockpit to be used immediately after installation.
Recommendation	As repository roles can change when a new version of the package is deployed, either do not use them directly but instead as a template for creating your own roles, or have a regular review process in place to verify that they still contain only privileges that are in line with your organization's security policy. Furthermore, if repository package privileges are granted by a role, we recommend that these privileges be restricted to your organization's packages rather than the complete repository. Therefore, 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.
How to Verify	To verify whether a user or role has a particular role, execute the following statement, for example: `SELECT * FROM EFFECTIVE_ROLE_GRANTEES WHERE ROLE_NAME ='sap.hana.xs.admin.roles::HTTPDestAdministrator';`
Related Alert	No

More Information	For a list of all roles delivered with each component, see *SAP HANA Security Reference Information* > *Components Delivered as SAP HANA Content* in the *SAP HANA Security Guide*.

User Parameter CLIENT

Default	The `CLIENT` user parameter can be used to authorize named users in SAP HANA. Only a user with the USER ADMIN system privilege can change the value of the `CLIENT` parameter already assigned to other users. However, at runtime, any user can assign an arbitrary value to the `CLIENT` parameter either by setting the corresponding session variable or passing the parameter via placeholder in a query. While this is the desired behavior for technical users that work with multiple clients such as SAP Business Warehouse, S/4 HANA, or SAP Business Suite, it is problematic in named user scenarios if the `CLIENT` parameter is used to authorize access to data and not only to perform data filtering.
Recommendation	Prevent named users from changing the `CLIENT` user parameter themselves but allow technical users to do so in their sessions and/or queries.

How to Verify	To verify that users are generally not permitted to change the `CLIENT` user parameter, ensure that the parameter `[authorization] secure_client_parameter` in the `global.ini` file is set to true:
	`SELECT * FROM "M_INIFILE_CONTENTS" WHERE KEY='SECURE_CLIENT_PARAMETER';`
	To verify that only permitted roles or users can change the CLIENT user parameter, execute the following statement:
	`SELECT * FROM EFFECTIVE_PRIVILEGE_GRANTEES WHERE OBJECT_TYPE = 'SYSTEMPRIVILEGE' AND PRIVILEGE = 'CLIENT PARAMETER ADMIN';`

Related Alert	No
More Information	See *SAP Note 2582162 (How to Restrict Use of the CLIENT Parameter)* and the section on authorization in the *SAP HANA Administration Guide*.

Related Information

Predefined Users
Deactivate the SYSTEM User
Password Policy
System Privileges
System Views for Verifying Users' Authorization
Predefined Database (Catalog) Roles
Predefined Repository Roles
Components Delivered as SAP HANA Content
Restrict Use of the CLIENT User Parameter

SAP Note 2582162
System Views for Verifying Users' Authorization

8.1.2.2 Recommendations for Network Configuration

Recommendations for integrating SAP HANA securely into your network environment.

General Recommendations

For general recommendations, please read the section on network security in the *SAP HANA Security Guide*.

Open Ports

Default	During installation, ports such as SQL 3<instance_no>15 and HTTP 80<instance_no> are opened by default.
Recommendation	Only ports that are needed for running your SAP HANA scenario should be open. For a list of required ports, see the *SAP HANA Administration Guide*.
How to Verify	Verify opened ports at operating system level using Linux commands such as `netcat` or `netstat`.
Related Alert	No
More Information	See the section on communication channel security in the *SAP HANA Security Guide* and the section on ports and connections in the *SAP HANA Administration Guide*.

Internal Host Name Resolution in Single-Host System

Default	SAP HANA services use IP addresses to communicate with each other. Host names are mapped to these IP addresses through internal host name resolution, a technique by which the use of specific and/or fast networks can be enforced and communication restricted to a specific network. In single-host systems, SAP HANA services listen on the loopback interface only (IP address 127.0.0.1). In `global.ini` files, the `[communication] listeninterface` is set to **`.local`**.
Recommendation	Do not change the default setting.

How to Verify	Using SAP HANA cockpit, check which ports are listening. This information is available in the *Network Security Information* app in the *SAP HANA Security Overview* catalog. The value of the *Listening On* field should be *Local Network*. Alternatively, execute the following SQL statement: `SELECT * FROM "PUBLIC" . "M_INIFILE_CONTENTS" WHERE SECTION = 'communication' AND KEY = 'listeninterface';`
Related Alert	No
More Information	See the section about ports and connections in the *SAP HANA Administration Guide*.

Internal Host Name Resolution in Multiple-Host System

Default	In a distributed scenario with multiple hosts, the network needs to be configured so that inter-service communication is operational throughout the entire landscape. The default configuration depends on how you installed your system.
Recommendation	Multiple-host systems can run with or without a separate network definition for inter-service communication. The recommended setting depends accordingly: • If a separate network **is configured** for internal communication, the parameter `[communication] listeninterface` should be set to **`.internal`**. In addition, you should add key-value pairs for the IP addresses of the network adapters used for SAP HANA internal communication in the `[communication] internal_hostname_resolution` section. • If a separate network **is not configured** for internal communication, the parameter `[communication] listeninterface` should be set to **`.global`**. This setting exposes internal SAP HANA service ports, so it is strongly recommended that you secure internal SAP HANA ports with an additional firewall. **i Note** Communication properties are in the default configuration change blacklist (`multidb.ini`). This means that they cannot initially be changed in tenant databases. They must be changed from the system database. If appropriate for your scenario, you can remove these properties from the change blacklist. SAP HANA deployment scenarios are described in the *SAP HANA Master Guide*. For more information about how to edit the change blacklist, see the *SAP HANA Administration Guide*.

How to Verify	Check which ports are listening using the SAP HANA cockpit.
	
	This information is available in the *Network Security Information* app in the *SAP HANA Security Overview* catalog. The value of the *Listening On* field should be *Global Network* or *Internal Network*.
	
	Alternatively, execute the following SQL statements:
	
	`SELECT * FROM "PUBLIC" . "M_INIFILE_CONTENTS" WHERE SECTION = 'communication' AND KEY = 'listeninterface';`
	
	`SELECT * FROM "PUBLIC" . "M_INIFILE_CONTENTS" WHERE SECTION = 'internal_hostname_resolution';`

Related Alert	86 (Internal communication is configured too openly)
More Information	See the section on internal hostname resolution in the *SAP HANA Administration Guide*.

Host Name Resolution in System Replication

Default	The parameter `[system_replication_communication] listeninterface` parameter is set to **`.global`**.
Recommendation	The recommended setting depends on whether or not a separate network is defined for internal communication: • If a separate internal network channel **is configured** for system replication, the parameter `[system_replication_communication] listeninterface` parameter should be **`.internal`**. You also need to add key-value pairs for the IP addresses of the network adapters for the system replication in the `[system_replication_hostname_resolution]` section. • If a separate network **is not configured** for system replication, the parameter `[system_replication_communication] listeninterface` should be set to **`.global`**. However, in this case, it is important to secure communication using TSL/SSL and/or to protect the SAP HANA landscape with a firewall. In the `[system_replication_hostname_resolution]` section, add entries for all hosts of neighboring sites (at a minimum) or all hosts of own site as well as for all hosts of neighboring sites. In addition, set the parameter `[system_replication_communication] allowed_sender` to restrict possible communication to specific hosts. The parameter value must contain a list of the foreign hosts that are part of the SAP HANA system replication landscape. **i Note** Communication properties are in the default configuration change blacklist (`multidb.ini`). This means that they cannot initially be changed in tenant databases. They must be changed from the system database. If appropriate for your scenario, you can remove these properties from the change blacklist. SAP HANA deployment scenarios are described in the *SAP HANA Master Guide*. For more information about how to edit the change blacklist, see the *SAP HANA Administration Guide*.

How to Verify	To check the value of the above parameters, execute the following statements:
	`SELECT * FROM "PUBLIC" . "M_INIFILE_CONTENTS" WHERE SECTION = 'system_replication_communication' AND KEY = 'listeninterface';`
	`SELECT * FROM "PUBLIC" . "M_INIFILE_CONTENTS" WHERE SECTION = 'system_replication_communication' AND KEY = 'internal_hostname_resolution';`
	`SELECT * FROM "PUBLIC". "M_INIFILE_CONTENTS"WHERE SECTION = 'system_replication_communication' AND KEY = 'allowed_sender';`

Related Alert	No
More Information	See the section on hostname resolution for system replication in the *SAP HANA Administration Guide*.

Related Information

Communication Channels
Network Security
Ports and Connections
Internal Host Name Resolution
Host Name Resolution for System Replication

8.1.2.3 Recommendations for Data Encryption

Recommendations for data encryption and encryption key management

Instance SSFS Master Key

Default	The instance secure store in the file system (SSFS) protects internal root keys in the file system. A unique master key is generated for the instance SSFS in every installation.
Recommendation	If you received your system pre-installed from a hardware or hosting partner, we recommend that you change the master key of the instance SSFS immediately after handover to ensure that it is not known outside of your organization.
How to Verify	Using SAP HANA cockpit, check the change date of the master key. This information is available in SAP HANA cockpit on the database overview page.
Related Alert	84 (Insecure instance SSF encryption configuration)

More Information	See the section on server-side data encryption in the *SAP HANA Security Guide* and the section on changing the SSFS master keys in the *SAP HANA Administration Guide*.

System PKI SSFS Master Key

Default	The system public key infrastructure (PKI) SSFS protects the X.509 certificate infrastructure that is used to secure internal TLS/SSL-based communication. A unique master key is generated for the system PKI SSFS in every installation.
Recommendation	If you received your system pre-installed from a hardware or hosting partner, we recommend that you change the master key of the instance SSFS immediately after handover to ensure that it is not known outside of your organization.
How to Verify	Check the change date of the master key in the SAP HANA cockpit. This information is available in the SAP HANA cockpit on *Security and User Management* view of the *Database Overview* page.
Related Alert	84 (Insecure instance SSF encryption configuration)
More Information	See the section on server-side data encryption in the *SAP HANA Security Guide* and the section on changing the SSFS master keys in the *SAP HANA Administration Guide*.

Root Encryption Keys

Default	SAP HANA features the following data encryption services: • Data volume encryption • Redo log encryption • Data and log backup encryption • An internal encryption service available to applications requiring data encryption Unique root keys are generated for all services in every database.
Recommendation	If you received your system pre-installed from a hardware or hosting partner, we recommend that you change all root keys immediately after handover to ensure that they are not known outside of your organization.
How to Verify	Query system view `ENCRYPTION_ROOT_KEYS`.
Related Alert	No
More Information	See the sections on server-side data encryption in the *SAP HANA Security Guide* and the *SAP HANA Administration Guide*.

Encryption Key of the SAP HANA Secure User Store (hdbuserstore)

Default	The secure user store (`hdbuserstore`) is a tool installed with the SAP HANA client. It is used to store SAP HANA connection information, including user passwords, securely on clients. Information contained in the SAP HANA secure user store is encrypted using a unique encryption key.
Recommendation	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.
How to Verify	You know the encryption has been changed if the file `SSFS_HDB.KEY` exists in the directory where the SAP HANA client is installed.
Related Alert	No
More Information	See the sections on `hdbuserstore` in the *SAP HANA Security Guide* and *SAP HANA Administration Guide*, as well as SAP Note 2210637.

Data and Log Volume Encryption

Default	Data and log volume encryption are not enabled
Recommendation	We recommend that you enable data and log volume encryption immediately after installation or handover from your hardware or hosting partner, and after you have changed the root encryption keys for both services.
How to Verify	Execute the following statement: `SELECT * FROM M_ENCRYPTION_OVERVIEW WHERE SCOPE='LOG' OR SCOPE = 'PERSISTENCE'`
Related Alert	No
More Information	See the section on data and log volume encryption in the *SAP HANA Security Guide* and the section on enabling encryption of data and log volumes in the *SAP HANA Administration Guide*.

Related Information

Server-Side Data Encryption Services
Change the SSFS Master Keys [page 233]
Change the Secure User Store Encryption Key
Changing Encryption Root Keys [page 238]
Secure User Store (hdbuserstore)
SAP Note 2210637

Data and Log Volume Encryption
Enabling Encryption of Data and Log Volumes [page 241]

8.1.2.4 Recommendations for File System and Operating System

Recommendations for secure operating system access and data storage in the file system

General Recommendation

Stay up to date on security recommendations available for your operating system and consider them in the context of your implementation scenario and security policy.

See also the following SAP Notes:

- SAP Note 1944799 (SUSE Linux Enterprise Server 11.x for SAP Applications)
- SAP Note 2009879 (Red Hat Enterprise Linux (RHEL) 6.x)

Operating System Users

Default	Only operating system (OS) users that are needed for operating SAP HANA exist on the SAP HANA system, that is: • `sapadm` (required to authenticate to SAP Host Agent) • `<sid>`adm (required by the SAP HANA database) • Dedicated OS users for every tenant database if the system is configured for high isolation **i Note** There may be additional OS users that were installed by the hardware vendor. Check with your vendor.
Recommendation	Ensure that no additional unnecessary users exist.
How to Verify	Refer to your operating system documentation
Related Alert	No
More Information	See the section on predefned users in the *SAP HANA Security Guide*.

OS File System Permissions

Default	The access permission of files exported to the SAP HANA server can be configured using the `[import_export] file_security` parameter in the `indexserver.ini` configuration file. The default permission set is 640 (`[import_export] file_security=medium`).
Recommendation	Do not change default access permission of exported files. In addition, ensure that only a limited number of database users have the system privilege `IMPORT` and `EXPORT`.

How to Verify	• You can verify the parameter setting by executing the command:
	`SELECT * FROM "PUBLIC" . "M_INIFILE_CONTENTS" WHERE SECTION = 'import_export' AND KEY = 'file_security';`
	• You can verify which users or roles have the `IMPORT` or `EXPORT` privilege by executing the statement:
	`SELECT * FROM EFFECTIVE_PRIVILEGE_GRANTEES WHERE (OBJECT_TYPE = 'SYSTEMPRIVILEGE') AND (PRIVILEGE = 'EXPORT' OR PRIVILEGE='IMPORT');`
	• You can verify the permissions of directories in the file system using the SAP HANA database lifecycle manager (HDBLCM) resident program with installation parameter `check_installation`.

Related Alert	No
More Information	See the section on checking the installation of an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) in the SAP HANA Administration Guide, as well as SAP Note 2252941.

OS Security Patches

Default	OS security patches are not installed by default
Recommendation	Install OS security patches for your operating system as soon as they become available. If a security patch impacts SAP HANA operation, SAP will publish an SAP Note where this fact is stated. It is up to you to decide whether to install such patches.
How to Verify	Refer to your operating system documentation
Related Alert	No
More Information	• SAP Note 1944799 (SUSE Linux Enterprise Server 11.x for SAP Applications) • SAP Note 2009879 (Red Hat Enterprise Linux (RHEL) 6.x)

OS sudo Configuration

Default	Users have to either specify the root password or be part of a dedicated user group to be able to run arbitrary commands as root.

Recommendation	Do not change your sudo configuration to allow users such as <sid>adm to use sudo to run arbitrary commands as root without specifying the root password.
How to Verify	Check the `/etc/sudoers` file. The specific configuration may vary with your Linux distribution, but configuration options to look for are: • `Defaults targetpw` This setting requires the root password to be provided when running sudo in general. • `ALL ALL=(ALL) ALL` This should only be used if `Defaults targetpw` is also set. If you use the storage connector option to mount SAP HANA volumes, during SAP HANA installation your sudo configuration is modified to allow <sid>adm to run a dedicated set of commands as root, such as: `<sid>adm ALL=NOPASSWD: /sbin/multipath,/sbin/multipathd,/etc/init.d/multipathd,/usr/bin/sg_persist,/bin/mount [...]` This is intentional and does not pose a security risk. However, <sid>adm should not be able to run arbitrary commands as root without proper authentication.
Related Alert	No
More Information	See the sudo and sudoers documentation (man 8 sudo, man 5 sudoers)

Related Information

Predefined Users
Check the Installation Using the Command-Line Interface
Secure User Store (hdbuserstore)
SAP Note 2252941
SAP Note 1944799
SAP Note 2009879

8.1.2.5 Recommendations for Auditing

Recommendations for audit configuration

Auditing

Default	Auditing is disabled by default.

Recommendation	Verify whether auditing is required by your security concept, for example to fulfill specific compliance and regulatory requirements.
How to Verify	Check the status of auditing in the SAP HANA cockpit This information is available on the *Auditing* card of the *SAP HANA Security Overview* catalog. Alternatively, you can execute the following statement: `SELECT * FROM "PUBLIC" . "M_INIFILE_CONTENTS" WHERE SECTION = 'auditing configuration' AND KEY = 'global_auditing_state';`
Related Alert	No
More Information	See the sections on auditing in the *SAP HANA Security Guide* and the *SAP HANA Administration Guide*.

Audit Trail Target: syslog

Default	The default audit trail target is syslog (`SYSLOGPROTOCOL`) for the system database
Recommendation	If you are using syslog, ensure that it is installed and configured according to your requirements (for example, for writing the audit trail to a remote server).
How to Verify	Refer to your operating system documentation
Related Alert	No
More Information	See the section on audit trails in the *SAP HANA Security Guide* and your operating system documentation.

Audit Trail Target: CSV Text File

Default	The audit trail target CSV text file (`CSVTEXTFILE`) is not configured by default
Recommendation	Do not configure CSV text file (`CSVTEXTFILE`) as an audit trail target in a production system as it has severe restrictions.
How to Verify	Check the configured audit trail targets in the *Auditing* of the SAP HANA cockpit Alternatively, execute the following statements: • `SELECT * FROM "PUBLIC" . "M_INIFILE_CONTENTS" WHERE SECTION = 'auditing configuration' AND VALUE = 'CSVTEXTFILE';` • `SELECT * FROM "PUBLIC"."AUDIT_POLICIES" WHERE TRAIL_TYPE='CSV';`
Related Alert	No
More Information	See the section on audit trails in the *SAP HANA Security Guide*.

Related Information

Auditing Activity in SAP HANA
Audit Trails
Auditing Activity in the SAP HANA Database
CREATE AUDIT POLICY Statement (Access Control)

8.1.2.6 Recommendations for Trace and Dump Files

Recommendations for handling trace and dump files

Trace Files

Default	Basic tracing of activity in database components is enabled by default, with each database service writing to its own trace file. Other traces (for example, SQL trace, expensive statements trace, performance trace) must be explicitly enabled. Users with the system privilege `CATALOG READ` can read the contents of trace files in the SAP HANA studio. At operating system level, any user in the `SAPSYS` group can access the trace directory: `/usr/sap/<SID>/HDB<instance>/<host>/trace/<db_name>`
Recommendation	• Enable tracing to troubleshoot specific problems only and then disable. • Exercise caution when setting or changing the trace level. A high trace level may expose certain security-relevant data (for example, database trace level `DEBUG` or SQL trace level `ALL_WITH_RESULTS`). • Delete trace files that are no longer needed.
How to Verify	• You can check which traces are enabled and how they are configured in the Administration editor of the SAP HANA studio on the *Trace Configuration* tab. • You can view trace files in the Administration editor of the SAP HANA studio on the *Diagnosis Files* tab and using the SAP HANA Database Explorer, which is integrated into the SAP HANA cockpit and SAP Web IDE for SAP HANA.
Related Alert	No
More Information	See the section on security risks of trace and dump files in the *SAP HANA Security Guide* and the section on configuring traces in the *SAP HANA Administration Guide*.

Dump Files

Default	The system generates core dump files (for example, crash dump files) automatically. Runtime (RTE) dump files can be triggered explicitly, for example by using the SAP HANA database management console (`hdbcons`) or as part of a full system information dump (`fullSystemInfoDump.py`). RTE dump files must be generated by the `<sid>`adm user. **⚠ Caution** Technical expertise is required to use `hdbcons`. To avoid incorrect usage, use `hdbcons` only with the guidance of SAP HANA development support. To create RTE dump files in a running system as part of a full system information dump in the SAP HANA studio, a user requires the `EXECUTE` privilege on procedure `SYS.FULL_SYSTEM_INFO_DUMP_CREATE`. Dump files are stored in the trace directory and have the same access permissions as other trace files (see above). Runtime dump files created as part of a full system information dump can be retrieved by users with the `EXECUTE` privilege on the procedure `SYS.FULL_SYSTEM_INFO_DUMP_RETRIEVE` using the SAP HANA studio. At operating system level, any user in the `SAPSYS` group can access their storage location: `/usr/sap/SID/SYS/global/sapcontrol/snapshots`
Recommendation	• Generate runtime dump files to analyze specific error situations only, typically at the request of SAP support. • Delete dump files that are no longer needed.
How to Verify	• You can view core dump files in the Administration editor of the SAP HANA studio on the *Diagnosis Files* tab. • You can download the file collections generated by a full system information dump in the Administration editor of the SAP HANA studio on the *Diagnosis Files* tab.
Related Alert	No
More Information	See the section on security risks of trace and dump files in the *SAP HANA Security Guide* and the section on collecting diagnosis information for SAP Support in the *SAP HANA Administration Guide*.

Related Information

Security Risks of Trace, Dump, and Captured Workload Files
Collecting Diagnosis Information for SAP Support

8.1.2.7 Recommendations for Tenant Database Management

Recommendations for securely configuring tenant databases

SAML-Based User Authentication

Default	All tenant databases use the same trust store as the system database for SAML-based user authentication
Recommendation	To prevent users of one tenant database being able to log on to other databases in the system (including the system database) using SAML, create individual certificate collections with the purpose **SAML** and **SSL** in every tenant database. In addition, specify a non-existent trust store for every tenant database using the `[communication] sslTrustStore` property in the `global.ini` file.
How to Verify	Execute the following statements: • In the tenant database: `SELECT * FROM PSES WHERE PURPOSE ='SAML' OR PURPOSE ='SSL';` • In the system database: `SELECT * FROM SYS_DATABASES.M_INIFILE_CONTENTS WHERE DATABASE_NAME='<TENANT_DB_NAME>' AND SECTION='communication' AND KEY = 'ssltruststore';`
Related Alert	No
More Information	See the sections on SSL configuration on the SAP HANA server and certficate collections in the *SAP HANA Security Guide*.

Configuration Blacklist

Default	A configuration change blacklist (`multidb.ini`) is delivered with a default configuration. The parameters contained in the blacklist can only be changed by a system administrator in the system database, not by the administrators of individual tenant databases.
Recommendation	Verify that the parameters included in the `multidb.ini` file meet your requirements and customize if necessary.
How to Verify	To see which parameters are blacklisted, execute the statement: `SELECT * FROM "PUBLIC". "M_INIFILE_CONTENTS" WHERE FILE_NAME = 'multidb.ini';`
Related Alert	No

More Information	See the section on default blacklisted system properties in tenant databases in the *SAP HANA Security Guide* and the section on how to prevent changes to system properties in tenant databases in the *SAP HANA Administration Guide*.

Restricted Features

Default	To safeguard and/or customize your system, it is possible to disable certain database features that provide direct access to the file system, the network, or other resources, for example import and export operations and backup functions. No features are disabled by default.
Recommendation	Review the list of features that can be disabled and disable those that are not required in your implementation scenario.
How to Verify	To see the status of features, query the system view `M_CUSTOMIZABLE_FUNCTIONALITIES`: `SELECT * FROM "PUBLIC". "M_CUSTOMIZABLE_FUNCTIONALITIES";`
Related Alert	No
More Information	See the section on restricted features in tenant databases in the *SAP HANA Security Guide* and the section on how to disable features on tenant databases in the *SAP HANA Administration Guide*.

Related Information

TLS/SSL Configuration on the SAP HANA Server
Certificate Collections [page 289]
Default Blacklisted System Properties in Tenant Databases [page 393]
Prevent Changes to System Properties in Tenant Databases
Restricted Features in Tenant Databases
Disable Features on a Tenant Database

8.1.3 Monitoring Critical Security Settings

SAP HANA has many configuration settings that allow you to customize your system for your implementation scenario and system environment. Some of these settings are important for the security of your system. Misconfiguration could leave your system vulnerable. The SAP HANA cockpit allows you to monitor several critical security settings at a glance.

> **i Note**
>
> In addition to using SAP HANA cockpit to monitor critical security settings, refer to *SAP HANA Security Checklists and Recommendations*. This document provides more detailed information as well as recommendations for many settings.

Related Information

View Status of Security Settings [page 222]
Overview of Security Options [page 223]
SAP HANA Security Checklists and Recommendations [page 197]

8.1.3.1 View Status of Security Settings

You can view the status of critical security settings of the SAP HANA database in the SAP HANA cockpit on the *Database Overview* page.

Prerequisites

- You have the authorization to see security-related information.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, review the status of the security settings.

   For more information, see *Overview of Security Options*.
2. To find more detailed information and functions, drill down.

Related Information

Overview of Security Options [page 223]

8.1.3.2 Overview of Security Options

The *Database Overview* offers information about important security settings, and links to further information and configuration options.

Cards

Card	Description	Required Authorization
Data Encryption	Indicates the status of data volume encryption, log volume encryption, and backup encryption The on/off switches allow you to enable or disable each type of encryption. **⚠ Caution** Do not enable data volume encryption in an existing operational database without having first read the section *Enabling and Disabling Encryption of Data and Log Volumes*. If you are connected to the system database, you will also see when the master keys of the secure stores in the file system (SSFS) were changed. This card opens the *Data Encryption* page, where you can see more information about the encryption status, enable or disable encryption, and change encryption keys.	System privilege CATALOG READ, ENCRYPTION ROOT KEY ADMIN (to enable/disable encryption), and RESOURCE ADMIN (to view SSFS master key information)

Card	Description	Required Authorization
Auditing	Indicates whether auditing is enabled in the database, the number of audit policies, and the configured audit trail target The *Auditing* switch allows you enable or disable auditing in the database. If a firefighter policy is active in the database (that is, a policy that audits all the actions of a particular user), this is also indicated. This card opens the *Auditing* page, where you can see more detailed information about audit policies, as well as create new ones. You can also make changes to global auditing settings.	To see information about auditing status audit policies, you need the system privileges AUDIT ADMIN or CATALOG READ. To see information about audit trail targets, you need the system privilege INIFILE ADMIN.
Authentication	Indicates the status of the password policy (default or customized), the user authentication mechanisms configured for single sign-on in the database, and when the password of the SYSTEM user was last changed This card opens the *Password Policy and Blacklist* page where you can see and edit the password policy and blacklist.	System privilege CATALOG READ To see the password blacklist on opening the *Password Policy and Blacklist* page, you need SELECT privilege on _SYS_SYS_PASSWORD_BLACKLIST (_SYS_SECURITY).

Anonymization Report

Link	Description	Authorization
View available anonymization views	Opens the *Anonymization Report* page, where you can view all calculation views with anonymization node views configured	CATALOG READ

Security Related Links

Link	Description	Authorization
Certificate Store	Opens the *Certificate Store* page, where you can import certificates into the certificate store	CATALOG READ, SSL ADMIN, USER ADMIN,TRUST ADMIN, or CERTIFICATE ADMIN
Certificate Collections	Opens the *Certificate Collections* page, where you can create and configure certificate collections	CATALOG READ, SSL ADMIN, USER ADMIN,TRUST ADMIN, or CERTIFICATE ADMIN
View network security information	Opens the *Network Security Information* page, where you can see more detailed information about network configuration.	CATALOG READ
SAML Identity Providers	Opens the *SAML Identity Providers* page, where you can see existing identity providers in the database and add new ones	USER ADMIN
JWT Identity Providers	Opens the *JWT Identity Providers* page, where you can see existing identity providers in the database and add new ones	USER ADMIN
Security Checklist	Opens the document *Security Checklist* on SAP Help Portal	No additional authorization required
Security Administration Help	Opens the SAP HANA documentation that describes those security administration tasks that you can perform using the SAP HANA cockpit	No additional authorization required
SAP HANA Security Website	Opens the SAP HANA security website	No additional authorization required

User & Role Management

Link	Description	Authorization
User Management	Opens the *User Management* page, where you view and manage database users	CATALOG READ to view users and USER ADMIN to create and manage users
Role Assignment	Opens the *Role Assignment* page, where you can see which roles are assigned to a user create and assign roles	CATALOG READ to view roles and ROLE ADMIN to assign roles

Link	Description	Authorization
Privilege Assignment	Opens the *Assign privileges to users* page, where you can see which roles are assigned to a user create and assign roles	CATALOG READ to view users; to assign privileges, USER ADMIN and the privileges required to grant specific privileges to the user
Role Management	Opens the *Role Management* page, where you view and manage roles in the database	CATALOG READ to view roles; to create and manage roles, ROLE ADMIN and the privileges required to grant specific privileges to roles
User Group Management	Opens the *User Group Management* page, where you view and manage user groups	CATALOG READ to view user groups; USER ADMIN to create user groups; object privilege USERGROUP OPERATOR on the group to edit user groups configured for exclusive administration
Authorization Dependency Viewer	You can use the Authorization Dependency Viewer to analyze database objects and their dependencies.	CATALOG READ

Related Information

Enabling Encryption of Data and Log Volumes [page 241]

8.1.3.3 Network Security Details

You can view important configuration settings related to secure internal SAP HANA communication and secure external SQL client communication on the *View network security communication* page of SAP HANA cockpit.

> **i** Note
>
> For more information about how to configure secure communication, see the *SAP HANA Security Guide*.

General Settings

Field	Description
Cryptographic Provider	The cryptographic service provider being used by the SAP HANA server
Maximum TLS/SSL Protocol Version Accepted	The maximum TLS/SSL protocol version accepted
Minimum TLS/SSL Protocol Version Accepted	The minimum TLS/SSL protocol version accepted

Field	Description
Allowed TLS/SSL Cipher Suites	The encryption algorithms allowed for TLS/SSL connections This value depends on the cryptographic service provider used. The default values are *PFS:HIGH::EC_HIGH:+EC_OPT* (CommonCryptoLib) and *ALL:!ADH:!LOW:!EXP:!NULL:@STRENGTH* (OpenSSL).

External JDBC/ODBC Communication

Field	Description
Enforce TLS/SSL for SQL Connections	Indicates whether all clients communicating with the SAP HANA database via the SQL interface are required to use a secured connection The database refuses SQL connection attempts that don't use TLS/SSL.
Key Store	The key store file that contains the server's private key(s)
Trust Store	The trust store file that contains the server's public certificate(s)
Validate Client Certificates	Indicates whether or not the certificate of the communication partner is validated

Internal Communication

Field	Description
TLS/SSL Secured	Indicates whether or not internal communication channels are secured using TLS/SSL The following values are possible: • *Disabled* (default) • *System PKI* • *Manual configuration* For more information about these values, see *Server-Side TLS/SSL Configuration Properties for Internal Communication* in the *SAP HANA Security Guide*.

Field	Description
Listening On	Indicates the listening interface for internal SAP HANA connections The following values are possible: • *Local network* 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. • *Global network* In multiple-host systems without a separate internal network, SAP HANA services listen on all available network interfaces. Connections from remote machines are possible. **⚠ Caution** 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. • *Internal network* 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. For more information, see *Configuring the Network for Multiple Hosts* and *Configuring SAP HANA Inter-Service Communication* in the *SAP HANA Administration Guide*.
Internal Host Name Resolution	The IP addresses of the network adapters used for SAP HANA internal communication This is relevant for multiple-host systems with a separate internal network (service communication: *Internal network*).
Key Store	The key store file that contains the server's private key(s)
Trust Store	The trust store file that contains the server's public certificate(s)
Validate Client Certificates	Indicates whether or not the certificate of the communication partner is validated

Related Information

Server-Side TLS/SSL Configuration Properties for Internal Communication
Configuring SAP HANA Inter-Service Communication

Configuring the Network for Multiple Hosts

8.1.4 Managing Server-Side Data Encryption

SAP HANA features a number of encryption services for encrypting data at rest, as well as an internal encryption service available to applications with data encryption requirements. SAP HANA uses the secure store in the file system (SSFS) functionality to all root encryption keys.

For more detailed information about the encryption services supported by SAP HANA, see the *SAP HANA Security Guide*.

Related Information

Encryption Configuration [page 229]
Change the SSFS Master Keys [page 233]
Set the Root Key Backup Password [page 236]
Back Up Root Keys [page 237]
Changing Encryption Root Keys [page 238]
Enabling Encryption of Data and Log Volumes [page 241]
Enable Encryption of Data and Log Backups [page 248]
Disable Data Encryption [page 250]
Import Backed-Up Root Keys [page 251]

8.1.4.1 Encryption Configuration

We recommend that you configure data encryption immediately after handover of your system from a hardware or hosting partner.

First-Time Configuration

The following figure shows the recommended process for configuring encryption in your SAP HANA system for the first time.

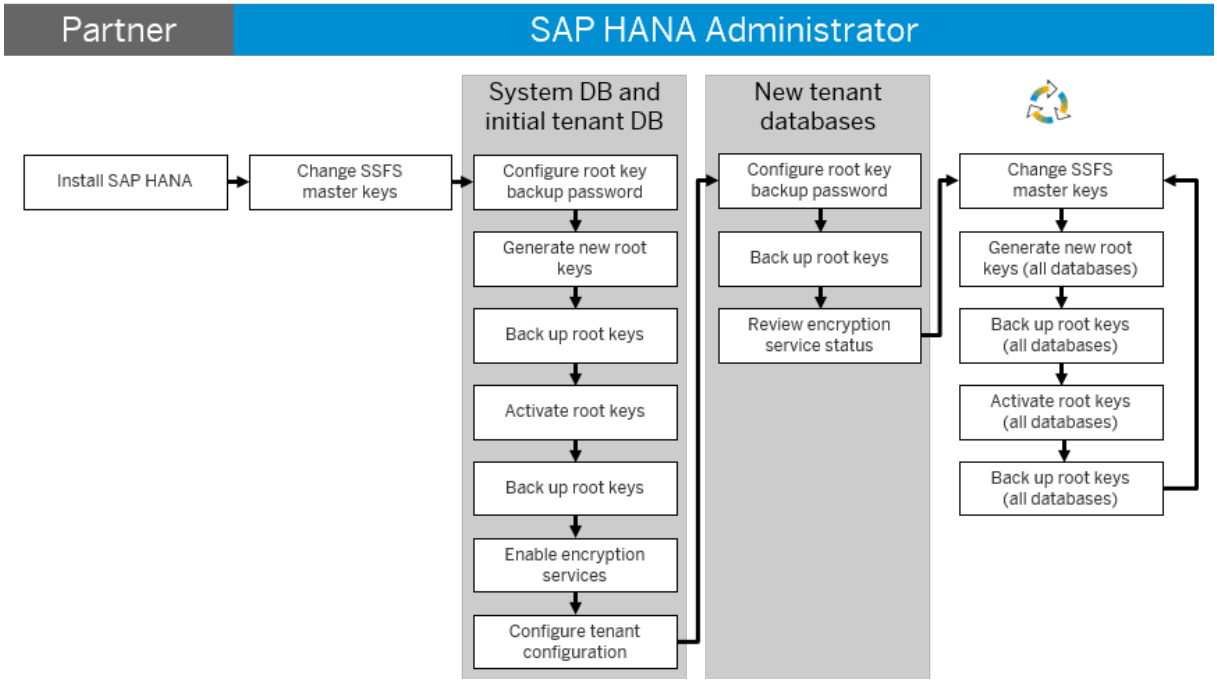

Immediately after system handover from your hardware or hosting partner, perform the following high-level steps.

Location	Steps	More Information
On the SAP HANA server	Change the master keys of the instance SSFS, if used, and the system PKI SSFS.(1)	Unique master keys are generated for the instance SSFS, if used, and the system PKI SSFS during installation or update. However, if you received your system pre-installed from a hardware or hosting partner, we recommend that you change them immediately after handover to ensure that they are not known outside of your organization. You can also change the master keys any time later.
In the system database	1. Set the password for the root key backup for the system database. This password is required to securely back up root keys and subsequently restore backed-up root keys during data recovery.(2)	The password is stored in the secure store along with the other root keys and used whenever you create a backup of the encryption root keys. The password should also be stored in a separate safe location. You will need to enter it to restore the secure store content before a database recovery. Losing this password may result in the database being unrecoverable.

Location	Steps	More Information
	2. Change the encryption root keys for all encryption services (data volume encryption, redo log encryption, data and log backup encryption, and internal application encryption) in the system database.(3) 1. Generate new root keys. 2. Back up all root keys to a root key backup file in a secure location. 3. Activate the new root keys. 4. Back up all root keys.	Unique root keys are generated during installation or database creation. However, if you received SAP HANA from a hardware or hosting partner, we recommend that you change them immediately after handover to ensure that they are not known outside of your organization. You can also change root keys any time later. You must back up all keys after you generate or activate a key of any type. This ensures that you always have an up-to-date backup of your root keys available for recovery. Store the root key backup file in a safe location. Losing this file may result in the database being unrecoverable.
	3. Enable the required encryption services: Data volume encryption, redo log encryption, and data and log backup encryption in the system database.(4)	Although SAP HANA provides you with the flexibility to encrypt data volumes, redo logs, and backups independently of each other, if you require full protection in the persistence layer, we recommend that you enable all services. It is not necessary to enable the internal application encryption service explicitly. It is available automatically to requesting applications.
	4. Configure how you want encryption to be handled in new tenant databases. You can do so with the following parameters in the `database_initial_encryption` section of the `global.ini` configuration file: • `persistence_encryption` (default: off) • `log_encryption` (default: off) • `backup_encryption` (default: off) • `encryption_config_control` (default: local_database)	By default, all encryption services are disabled in new tenant databases and only tenant database administrators can enable them.
In the first tenant database (if automatically created during installation)	1. Set the password for the root key backup for the first tenant database.(2)	The password is required to securely back up root keys and subsequently restore backed-up root keys during data recovery.

Location	Steps	More Information
	2. Change the encryption root keys for all encryption services in the first tenant database.(3)	See step 2 above (system database).
	3. Enable the required encryption services in the first tenant database.(4)	By default, only the tenant database administrator can do this in the tenant database. See *Encryption Configuration Control*.
In subsequent tenant databases	1, Set the password for the root key backup for the tenant database.(2)	The password is required to securely back up root keys and subsequently restore backed-up root keys during data recovery.
	2. Back up all root keys to a root key backup file in a secure location.	It is not necessary to change the root keys in new tenant databases. Unique root keys are generated on database creation and cannot be known outside of your organization.
	3. Change the status of encryption services in the tenant database if required.(4) Encryption services are configured in line with the parameters in the `database_initial_encryption` section of the `global.ini` configuration file as described above.	Who can enable or disable encryption services depends on how the parameter `encryption_config_control` is configured. See *Encryption Configuration Control*.
During operation	Periodically change the master keys of the instance SSFS, if used, and the system PKI SSFS in line with your security policy. Periodically change the encryption root keys in all databases in line with your security policy.	

(1) If the instance SSFS is used in a system-replication configuration, you change the instance SSFS master key on the primary system. To trigger replication of the new key to the secondary system, you must subsequently restart the secondary system. In multi-tier system replication scenarios involving three systems, restart the tier-2 secondary system first, then the tier-3 secondary system. If a secondary system takes over from its replication source before the new master key has been replicated, all systems registered will use the old key from the former secondary system instead.

(2) In a system-replication configuration, set the root key backup password in the primary system only. The password will be propagated to all secondary systems. The secondary systems must be running and replicating.

(3) In a system-replication configuration, change root keys in the primary system only. New keys will be propagated to all secondary systems. The secondary systems must be running and replicating.

(4) In a system-replication configuration, enable (or disable) encryption in the primary system only. The setting will be propagated to all secondary systems. The secondary systems must be running and replicating.

Configuration After Update from a Single-Container System

If you updated from a single-container system, your system has a system database and one tenant database. The existing data encryption configuration is retained. Note the following:

- The SSFS master keys for the system remain unchanged and the instance SSFS is still used.
- Existing encryption root keys are the encryption root keys of the tenant database. The update process generates new unique root keys for the system database.
- If a root key backup password existed before update, it is the root key backup password of the tenant database. The system database will not have a root key backup password set. You must set this password manually after the update.
- Encryption services that were enabled before update are enabled in both the system database and the tenant database.

Related Information

Change the SSFS Master Keys [page 233]
Set the Root Key Backup Password [page 236]
Changing Encryption Root Keys [page 238]
Enable Encryption [page 242]

8.1.4.2 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.

> **i** Note
>
> If you have replaced the instance SSFS with the local secure store (LSS), you still have a master key for the system PKI SSFS.

Prerequisites

- You have shut down the SAP HANA system.
- You have the credentials of the operating system user (`<sid>`adm) that was created when the system was installed.

Context

You can change the SSFS master keys using the command line tool `rsecssfx`, which is installed with SAP HANA and available at `/usr/sap/<SID>/HDB<instance>/exe`.

Before changing the SSFS master keys, note the following:

- In a distributed SAP HANA system, every host must be able to access the file location of the instance SSFS master key.
- The SSFS master keys only have to be changed once for the whole instance and not per tenant database.
- If the instance SSFS is used in a system-replication configuration, you change the instance SSFS master key on the primary system. To trigger replication of the new key to the secondary system, you must subsequently restart the secondary system. In multi-tier system replication scenarios involving three systems, restart the tier-2 secondary system first, then the tier-3 secondary system. If a secondary system takes over from its replication source before the new master key has been replicated, all systems registered will use the old key from the former secondary system instead.

Procedure

1. Log on to the SAP HANA system host as the operating system user, `<sid>`adm.
2. Change the master key of the **instance SSFS**, if used, 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)
   ```

   For script languages bash and csh the syntax is:

   ```
   rsecssfx changekey `rsecssfx generatekey -getPlainValueToConsole`
   ```

   > **i** Note
   >
   > The command uses backticks.

   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 `/usr/sap/<sid>/SYS/global/hdb/security/ssfs`. If you change the default path, you may need to reconfigure it in the event of a system rename.

3. Re-encrypt the **system PKI SSFS** with a new key with the following command:

   ```
   export RSEC_SSFS_DATAPATH=/usr/sap/<SID>/SYS/global/security/rsecssfs/data
   ```

```
export RSEC_SSFS_KEYPATH=<path to current key file>
rsecssfx changekey $(rsecssfx generatekey -getPlainValueToConsole)
```

> **i Note**
>
> The default path of the key file is `/usr/sap/<sid>/SYS/global/security/rsecssfs/key`. If you change the default path, you may need to reconfigure it in the event of a system rename.

For script languages bash and csh the syntax is:

```
rsecssfx changekey `rsecssfx generatekey -getPlainValueToConsole`
```

> **i Note**
>
> The command uses backticks.

4. Restart the SAP HANA system.

Next Steps

In a system-replication setup, perform the following steps:

1. Configure the location of the instance SSFS master key file on the secondary system(s). The file itself will be automatically copied when you restart the secondary system(s)
2. Restart the secondary system(s) to trigger the replication of the key files.

> **→ Remember**
>
> In multi-tier system replication scenarios involving three systems, restart the tier-2 secondary system first, then the tier-3 secondary system.

For file system-based copies of SAP HANA database installations, you must manually save and restore the instance SSFS master key file. Otherwise data loss can occur.

In regular backup and recovery scenarios, the SSFS must always be restored from a recent root key backup before a database recovery, unless:

- You have never changed the redo log encryption key.
- You are performing a recovery into the same database from which the backup was taken, and the database's secure store is intact and contains the latest root key changes.

> **i Note**
>
> It is not necessary to save the system PKI SSFS key file. The system will generate a new system PKI SSFS automatically if required.

Related Information

Stop a System

Configuring SAP HANA System Properties (INI Files)
Server-Side Data Encryption Services
Import Backed-Up Root Keys [page 251]
SAP Note 2183624

8.1.4.3 Set the Root Key Backup Password

The root key backup password is required to securely back up the root keys of the database and subsequently to restore the backed-up root keys during data recovery.

Prerequisites

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

  .

Procedure

1. From the *Database Overview*, with the *Security and User Management* or *All* view selected, navigate to the *Data Encryption* area.
2. Open the *Data Encryption Configuration* page by clicking the *Data Encryption* card.
3. Choose *Manage Keys*.
4. On the *Manage Keys* page, click *Set Root Key Backup Password* and specify the password.

   The length and layout of the password must be in line with the database's password policy.

   > ⚠ Caution
   >
   > If the root key backup already has a password, it will be overwritten.

   > i Note
   >
   > In a system-replication configuration, set the root key backup password in the primary system only. The password will be propagated to all secondary systems. The secondary systems must be running and replicating.

Results

The password is set and stored in the secure store in the file system (SSFS) together with other root keys. The password must be set to enable root keys to be backed up securely. You must provide this password to import

root keys from the backup into the database before starting a recovery. All root key backups taken after the password is set use this password to protect the backup files.

For more information about root key backup, see the *SAP HANA Security Guide*.

> ⚠ Caution
>
> The password is stored in the secure store along with the other root keys and used whenever you create a backup of the encryption root keys. The password should also be stored in a separate safe location. You will need to enter it to restore the secure store content before a database recovery. Losing this password may result in the database being unrecoverable.

> → Tip
>
> To verify that a password is the same as the one stored in the instance SSFS, use the statement `ALTER SYSTEM VALIDATE ENCRYPTION ROOT KEYS BACKUP PASSWORD`.

Related Information

Root Key Backup
Password Policy Details [page 275]

8.1.4.4 Back Up Root Keys

After you have generated or activated new encryption root keys, or created a new tenant database with new root keys, you must back up all root keys.

Prerequisites

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- You have set the root key backup password.
- The external location to which you plan to back up root keys is accessible.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. From the *Database Overview*, with the *Security and User Management* or *All* view selected, navigate to the *Data Encryption* area.

   The Data Encryption Configuration page appears.

2. Choose *Manage Keys*.
3. Choose *Back Up Root Keys*.
4. Save the root key backup file to a secure location.

   > ⚠ Caution
   >
   > Store the root key backup file in a safe location. If this file is lost, it may not be possible to recover the database.

5. Optional: To ensure that the backup file can be recovered, validate the password for the root key backup file. Log on to the SAP HANA server as operating system user <sid>adm, and use the following command in the `hdbnsutil` tool:

   ```
   cd /usr/sap/<sid>/<HDBinstance_no>/exe
   ./hdbnsutil -validateRootKeysBackup <filename> [--password=<passphrase>]
   ```

   > → Recommendation
   >
   > We recommend that you do not enter the password on the command line. You will be interactively prompted to enter it. In this way, you avoid unintentionally leaving the password in the command history and making it visible in process monitoring tools provided by the operating system.

Related Information

Set the Root Key Backup Password [page 236]
ALTER SYSTEM SET ENCRYPTION ROOT KEYS BACKUP PASSWORD Statement (System Management)
ALTER SYSTEM VALIDATE ENCRYPTION ROOT KEYS BACKUP PASSWORD Statement (System Management)
ENCRYPTION_ROOT_KEYS_EXTRACT_KEYS Function (Security)
Back Up Root Keys (hdbnsutil)

8.1.4.5 Changing Encryption Root Keys

Unique root keys are generated during installation or database creation. However, if you received SAP HANA from a hardware or hosting partner, we recommend that you change them immediately after handover to ensure that they are not known outside of your organization. You can also change root keys any time later.

Change the root keys for the following encryption services immediately after handover of your system and periodically during operation:

- Data volume encryption
- Redo log encryption
- Data and log backup encryption
- Internal application encryption

It is important to always change encryption root keys as follows:

1. Generate new root keys.

2. Back up all root keys.
3. Activate new root keys.
4. Back up all root keys.

You must back up all keys after you generate or activate a key of any type. This ensures that you always have an up-to-date backup of your root keys available for recovery.

> **i** Note
>
> In a system-replication configuration, change root keys in the primary system only. New keys will be propagated to all secondary systems. The secondary systems must be running and replicating.

You can change root keys using the SAP HANA cockpit or from the command line.

Related Information

Encryption Key Management

8.1.4.5.1 Change Root Keys

The process for changing encryption root keys involves first generating the new keys and backing them up, and then activating them and backing them up again. You can change all root keys following this process on the *Manage Keys* page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- The external location to which you plan to back up root keys is accessible.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. From the *Database Overview*, with the *Security and User Management* or *All* view selected, navigate to the *Data Encryption* area.
2. Open the *Data Encryption* page by clicking the *Data Encryption* card, and then choose *Manage Keys*.
3. On the *Manage Keys* page, choose *Change Root Keys*.
4. If you have not already done so, set the root key backup password.

   The length and layout of the password must be in line with the database's password policy.

> ⚠ Caution
>
> The password is stored in the secure store along with the other root keys and used whenever you create a backup of the encryption root keys. The password should also be stored in a separate safe location. You will need to enter it to restore the secure store content before a database recovery. Losing this password may result in the database being unrecoverable.

5. Select the root keys that you want to change.
6. Back up all root keys to a secure location.

> ⚠ Caution
>
> Store the root key backup file in a safe location. Losing this file may result in the database being unrecoverable.

7. Activate the new keys.
8. Back up all root keys again.
9. Optional: To ensure that the backup file can be recovered, validate the password for the root key backup file. Log on to the SAP HANA server as operating system user <sid>adm, and use the following command in the `hdbnsutil` tool:

```
cd /usr/sap/<sid>/<HDBinstance_no>/exe
./hdbnsutil -validateRootKeysBackup <filename> [--password=<passphrase>]
```

> → Recommendation
>
> We recommend that you do not enter the password on the command line. You will be interactively prompted to enter it. In this way, you avoid unintentionally leaving the password in the command history and making it visible in process monitoring tools provided by the operating system.

Results

If encryption is enabled, new data is encrypted with the new root keys.

On the *Data Encryption* page, if the instance SSFS is used, the active version of changed root keys increments by one and the last changed date is updated. If LSS is used, the version shown is always -1.

Related Information

Set the Root Key Backup Password
Password Policy Details [page 275]
Password Policy Configuration Options

8.1.4.6 Enabling Encryption of Data and Log Volumes

You can enable data volume encryption and redo log encryption in a new SAP HANA database or in an existing operational database.

> → Recommendation
>
> Although SAP HANA provides you with the flexibility to encrypt data volumes, redo logs, and backups independently of each other, if you require full protection in the persistence layer, we recommend that you enable all services.

Related Information

Enable Data and Log Volume Encryption in a New SAP HANA Database [page 241]
Enable Encryption [page 242]

8.1.4.6.1 Enable Data and Log Volume Encryption in a New SAP HANA Database

The recommended time to enable data and log volume encryption is immediately after tenant database creation. If you received SAP HANA from a hardware or hosting partner, enable encryption after handover.

Prerequisites

- The tenant database has control of enabling encryption.
  By default, encryption can be enabled or disabled only directly in the tenant database and not from the system database.
  To see how your database is configured, query the system view SYS.M_ENCRYPTION_OVERVIEW. For more information about how to switch control, see *Encryption Configuration*.
  If the system database has control, encryption can only be enabled or disabled using SQL from the system database. For more information, see the *SAP HANA Administration Guide*.
- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- If necessary, you have changed and backed up the encryption root keys.
  SAP HANA generates unique root keys on installation or database creation. However, if you received SAP HANA from a hardware or hosting partner, we recommend that you change the root keys used for data volume encryption and redo log encryption. This ensures that the root keys are not known outside your organization. For more information, see the section on changing root encryption keys.

> **Note**
>
> In a system-replication configuration, change the root keys used for data volume encryption and log volume encryption in the primary system only. The new keys will be propagated to all secondary systems.

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. From the *Database Overview*, with the *Security and User Management* or *All* view selected, navigate to the *Data Encryption* area.
2. On the *Data Encryption* card, enable data volume encryption and redo log encryption with the on/off switch.

> **Note**
>
> You can also enable encryption on the *Data Encryption Configuration* page.

Results

All data persisted to data volumes is encrypted and all future redo log entries persisted to log volumes are encrypted.

You can monitor the progress of data volume encryption service by service on the *Data Encryption Configuration* page. Once encryption of a data volume has completed, the status changes to *Encrypted*.

Related Information

Encryption Configuration [page 229]
Changing Encryption Root Keys [page 238]
M_ENCRYPTION_OVERVIEW System View

8.1.4.6.2 Enable Encryption

You can enable data volume encryption, redo log encryption, and encryption of data and log backups in a new SAP HANA database or in an existing operational database.

Prerequisites

- To enable encryption for a tenant database, you know whether encryption configuration is controlled by the tenant database or the system database.
  - If the tenant database controls encryption configuration, encryption can only be enabled or disabled directly in the tenant database and not from the system database.
  - If the system database controls encryption configuration, encryption can only be enabled or disabled using SQL from the system database, with a user that has the system privilege DATABASE ADMIN.
- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- If necessary, you have changed and backed up the encryption root keys.

Context

It is recommended that you enable encryption in the system database and the tenant databases when they are created. In this way, you ensure that all the pages are encrypted. If you received SAP HANA from a hardware or hosting partner, you should enable encryption after handover and before importing your sensitive data.

If you enable encryption in an operational database, only the pages in use in the data volumes are encrypted. Pages in data volumes that are not in use may still contain old content, and are only overwritten and encrypted over time. This means that your data in data volumes will only be fully encrypted after some delay. In addition, only redo log entries that are created after encryption is enabled are encrypted. Redo log files that were created before encryption was enabled are not encrypted. Although encryption can be switched on at any point in time, unencrypted data can remain on disk. If this is not wanted, you will need to install a new database on a fresh hard drive, activate encryption, import a backup, and low-level erase the old disks.

You can enable encryption of full data backups, delta data backups, and log backups in the database at any time.

> → Recommendation
>
> Although SAP HANA provides you with the flexibility to encrypt data volumes, redo logs, and backups independently of each other, if you require full protection in the persistence layer, we recommend that you enable all services.

Procedure

> i Note
>
> In a system-replication configuration, enable (or disable) encryption in the primary system only. The setting will be propagated to all secondary systems. The secondary systems must be running and replicating.

Enable the required encryption service.

You can do this using SQL or SAP HANA cockpit.

Option	Description
SQL	If the tenant database has control: ○ Data volume encryption: `ALTER SYSTEM PERSISTENCE ENCRYPTION ON` ○ Redo log encryption: `ALTER SYSTEM LOG ENCRYPTION ON` ○ Backup encryption: `ALTER SYSTEM BACKUP ENCRYPTION ON` If the system database has control: ○ Data volume encryption: `ALTER DATABASE <database_name> PERSISTENCE ENCRYPTION ON` ○ Redo log encryption: `ALTER DATABASE <database_name> LOG ENCRYPTION ON` ○ Backup encryption: `ALTER DATABASE <database_name> BACKUP ENCRYPTION ON`
SAP HANA cockpit	Tenant database control only: 1. From the *Database Overview*, with the *Security and User Management* or *All* view selected, navigate to the *Data Encryption* area. 2. On the *Data Encryption* card, enable the relevant encryption service using the on/off switch. **i Note** You can also enable each encryption service on the *Data Encryption Configuration* page.

Results

Data volume encryption and redo log encryption

All data persisted to data volumes is encrypted and all future redo log entries persisted to log volumes are encrypted.

You can verify the status of data volume encryption and redo log encryption in the system view M_ENCRYPTION_OVERVIEW or on the *Data Encryption* tile of the SAP HANA cockpit.

Backup encryption

Backup encryption is enabled. Subsequent log backups, as well as full backups and delta data backups will be encrypted.

> **i Note**
>
> If backup encryption is active, a data snapshot is **not automatically encrypted**. For more information, see *SAP HANA Backup Encryption*.

Related Information

SAP Note 2159014
Encryption Configuration [page 229]
Changing Encryption Root Keys [page 238]
M_ENCRYPTION_OVERVIEW System View
ALTER SYSTEM PERSISTENCE ENCRYPTION Statement (System Management)
ALTER SYSTEM LOG ENCRYPTION Statement (System Management)
ALTER DATABASE Statement (Tenant Database Management)
Backup Encryption
SAP HANA Backup Encryption
ALTER SYSTEM BACKUP ENCRYPTION Statement (System Management)
Creating Backups

8.1.4.6.2.1 Enable Data and Log Volume Encryption with Database Re-Creation

The recommended way to enable data volume encryption and redo log encryption in an existing operational SAP HANA database is after first dropping and re-creating the tenant database.

Prerequisites

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- You have the system privilege `DATABASE ADMIN`.
- You have the privileges required to perform backup and recovery.

Context

Enabling data volume encryption and redo log encryption after re-creating your tenant database ensures that new encryption root keys are generated. In addition, it provides complete protection. If you enable encryption without re-creating the database, only the pages in use within the data volumes will be encrypted. Pages in data volumes that are not in use may still contain old content and will only be overwritten and encrypted over time. This means that your data in data volumes will only be fully protected after some delay. In addition, only future redo log entries will be encrypted. Existing redo log files are not encrypted.

For more information about this recommendation, see SAP Note 2159014.

The following is the overall process for enabling encryption with database re-creation. For more information on individual steps, see the corresponding section.

Procedure

1. Perform a data backup.
2. Drop the tenant database.
3. Clean the disk space.
4. Create the tenant database again.
5. Set the password for the root key backup in the tenant database.
6. Back up all root keys to a root key backup file (*.rkb) in a secure location.
7. Enable data volume encryption and redo log encryption.

   By default, encryption is initially disabled in a new database and can be enabled only in the tenant database. This initial configuration is controlled by the parameters in the `database_initial_encryption` section of the `global.ini` configuration file. To see how the database is configured, query the system view SYS.M_ENCRYPTION_OVERVIEW, or from the system database SYS_DATABASES.M_ENCRYPTION_OVERVIEW.
8. Recover your tenant database.

Results

All data persisted to data volumes is encrypted and all future redo log entries persisted to log volumes are encrypted.

Related Information

Encryption Configuration [page 229]
Create Data Backups [page 603]
Delete a Tenant Database [page 387]
Create a Tenant Database [page 381]
Set the Root Key Backup Password [page 236]
Back Up Root Keys [page 237]
Enable Data and Log Volume Encryption in a New SAP HANA Database [page 241]
Recover a Database [page 628]
M_ENCRYPTION_OVERVIEW System View
SAP Note 2159014

8.1.4.6.2.2 Enable Data and Log Volume Encryption Without Database Re-Creation

If it is not possible to drop and re-create your SAP HANA database to enable encryption and redo log encryption, for example, because it would result in too much downtime, you can enable encryption

immediately. However, this is not recommended because your data will only be fully protected after some delay.

Prerequisites

- The tenant database has control of enabling encryption.
  By default, encryption can be enabled or disabled only directly in the tenant database and not from the system database.
  To see how your database is configured, query the system view SYS.M_ENCRYPTION_OVERVIEW. For more information about how to switch control, see *Encryption Configuration*.
  If the system database has control, encryption can only be enabled or disabled using SQL from the system database. For more information, see the *SAP HANA Administration Guide*.
- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- If necessary, you have changed and backed up the encryption root keys.
  SAP HANA generates unique root keys on installation or database creation. However, if you received SAP HANA from a hardware or hosting partner, we recommend that you change the root keys used for data volume encryption and redo log encryption. This ensures that the root keys are not known outside your organization.

  > **i** Note
  >
  > In a system-replication configuration, change the root keys used for data volume encryption and log volume encryption in the primary system only. The new keys will be propagated to all secondary systems.

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

For maximum protection, we recommend that you drop and re-create your SAP HANA database before enabling data volume encryption and redo log encryption. If you enable encryption once the database has been operational, only the pages in use within the data volumes will be encrypted. Pages in data volumes that are not in use may still contain old content and will only be overwritten and encrypted over time. This means that your data in data volumes will only be fully protected after some delay. In addition, only future redo log entries will be encrypted. Existing redo log files are not encrypted.

For more information about this recommendation, see SAP Note 2159014.

Procedure

1. From the *Database Overview*, with the *Security and User Management* or *All* view selected, navigate to the *Data Encryption* area.

2. On the *Data Encryption* card, enable data volume encryption and redo log encryption with the on/off switch.

> **i Note**
>
> You can also enable encryption on the *Data Encryption Configuration* page.

Results

Encryption is now active for all new data saved to disk as of the next savepoint operation. Existing data starts being encrypted in the background. Only after this process has completed is all your data encrypted.

You can monitor the progress of data volume encryption service by service on the *Data Encryption Configuration* page. Once encryption of a data volume has completed, the status changes to *Encrypted*

> **→ Remember**
>
> Due to the shadow memory nature of SAP HANA database persistence, the data area may still contain outdated, unencrypted versions of pages.

All future redo log entries persisted to log volumes are encrypted. Previous unencrypted redo log entries remain unencrypted.

> **→ Remember**
>
> Existing redo log files will not be encrypted until they are overwritten.

Related Information

Encryption Configuration [page 229]
M_ENCRYPTION_OVERVIEW System View
SAP Note 2159014

8.1.4.7 Enable Encryption of Data and Log Backups

You can enable encryption of full data backups, delta data backups, and log backups in an SAP HANA database at any time.

Prerequisites

- The tenant database has control of enabling encryption.

By default, encryption can be enabled or disabled only directly in the tenant database and not from the system database.
To see how your database is configured, query the system view SYS.M_ENCRYPTION_OVERVIEW. For more information about how to switch control, see *Encryption Configuration*.
If the system database has control, encryption can only be enabled or disabled using SQL from the system database. For more information, see the *SAP HANA Administration Guide*.

- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- If necessary, you have changed and backed up the encryption root keys.
  SAP HANA generates unique root keys on installation or database creation. However, if you received SAP HANA from a hardware or hosting partner, consider changing the backup encryption root key to ensure it is not known outside your organization.

  > **i Note**
  >
  > In a system-replication configuration, change the root keys used for data volume encryption and log volume encryption in the primary system only. The new keys will be propagated to all secondary systems.

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. From the *Database Overview*, with the *Security and User Management* or *All* view selected, navigate to the *Data Encryption* area.
2. On the *Data Storage Security* card, enable backup encryption with the on/off switch.

   > **i Note**
   >
   > You can also enable backup volume encryption on the *Data Encryption Configuration* page.

Results

Backup encryption is enabled. Subsequent log backups, as well as full backups and delta data backups will be encrypted.

> **i Note**
>
> If backup encryption is active, a data snapshot is **not automatically encrypted**. Data in storage snapshots is encrypted as part of **data volume encryption**.

Related Information

Encryption Configuration [page 229]

Changing Encryption Root Keys [page 238]
SAP HANA Backup Encryption
Enabling Encryption of Data and Log Volumes [page 241]

8.1.4.8 Disable Data Encryption

Disabling data volume encryption triggers the decryption of all encrypted data. Newly persisted data is not encrypted. Disabling redo log encryption makes sure that future redo log entries are not encrypted when they are written to disk.

Prerequisites

- The tenant database has control of enabling encryption.
  By default, encryption can be enabled or disabled only directly in the tenant database and not from the system database.
  To see how your database is configured, query the system view SYS.M_ENCRYPTION_OVERVIEW. For more information about how to switch control, see *Encryption Configuration*.
  If the system database has control, encryption can only be enabled or disabled using SQL from the system database. For more information, see the *SAP HANA Administration Guide*.
- You have the system privilege `ENCRYPTION ROOT KEY ADMIN`.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. From the *Database Overview*, with the *Security and User Management* or *All* view selected, navigate to the *Data Encryption* area.
2. On the *Data Encryption* card, disable the relevant encryption service using the on/off switch.

> **i Note**
>
> You can also disable each encryption service on the *Data Encryption Configuration* page.

Results

Data volume encryption

Data starts being decrypted in the background. Depending on the size of the SAP HANA database, this process can be very time consuming. Only after this process has completed is all your data decrypted. Newly persisted data is not encrypted.

You can monitor the progress of data volume decryption service by service. Once decryption of a data volume has completed, the status changes to *Unencrypted*.

Redo log encryption

New redo log entries are not encrypted. Existing redo log entries are not decrypted. Log entries will only be fully unencrypted when all encrypted entries have been overwritten.

Backup encryption

New data backups, delta backups, and log backups are not encrypted. On an unencrypted data volume, data snapshots are also unencrypted.

Related Information

Encryption Configuration [page 229]
M_ENCRYPTION_OVERVIEW System View

8.1.4.9 Import Backed-Up Root Keys

Before performing a recovery from encrypted data and log backups, you must import the backed-up root keys. The imported keys are used to initialize the secure store. When the instance SSFS is used as the secure store, the set of keys from the root key backup file replaces the set of keys in the instance SSFS. When LSS is used as the secure store, the key import ensures that only the missing keys are added to the secure store; no keys are deleted.

Prerequisites

- You have the credentials of the operating system user (`<sid>`adm).
- You can log on to the system database and have the system privilege `DATABASE STOP`.
- If using the SQL statement RECOVER ENCRYPTION ROOT KEYS, you have the system privilege `DATABASE RECOVERY OPERATOR` or `DATABASE ADMIN`.
- The location of the root key backup file is accessible.
- You know the ID of the database whose root keys you want to import. You can determine the IDs of all tenant databases by executing the following SQL command in the system database:

```
SELECT DATABASE_NAME,
    CASE WHEN (DBID = '' AND
    DATABASE_NAME = 'SYSTEMDB')
    THEN 1
    WHEN (DBID = '' AND
    DATABASE_NAME <> 'SYSTEMDB')
    THEN 3
    ELSE TO_INT(DBID)
    END DATABASE_ID
FROM (SELECT DISTINCT DATABASE_NAME, SUBSTR_AFTER (SUBPATH,'.') AS DBID FROM
SYS_DATABASES.M_VOLUMES);
```

Context

The following options are available for importing root keys:

Option	Execute...	Database Mode	Scope
hdbnsutil tool	On the SAP HANA server	The tenant database must be offline.	The encryption root keys can be recovered for tenant databases and system databases.
SQL statement RECOVER ENCRYPTION ROOT KEYS	In the system database	The tenant database must be offline.	The encryption root keys can only be recovered for tenant databases.

Procedure

1. Log on to the SAP HANA server as operating system user `<sid>adm`.
2. Validate that you have the password for the root key backup file:

```
cd /usr/sap/<sid>/<HDBinstance_no>/exe
./hdbnsutil -validateRootKeysBackup <filename> [--password=<passphrase>]
```

> → Recommendation
>
> We recommend that you do not enter the password on the command line (`--password`). By omitting this parameter, you will be interactively prompted to enter it. In this way, you avoid unintentionally leaving the password in the command history and making it visible in process monitoring tools provided by the operating system.

3. In the system database, stop the tenant database to be recovered.

   You can do this in the SAP HANA cockpit or by executing the statement `ALTER SYSTEM STOP <database_name>`.

4. Import the backed-up root keys using either the `hdbnsutil` program or the SQL statement RECOVER ENCRYPTION ROOT KEYS:

   - Using the `hdbnsutil` program

     Execute the following command:

```
cd /usr/sap/<sid>/<HDBinstance_no>/exe
./hdbnsutil -recoverRootKeys <filename> --dbid=<dbid> --
password=<passphrase> --type=ALL
```

     - `<dbid>` is the tenant database ID.
     - The `<type>` option is the root key type and also accepts the values `PERSISTENCE`, `LOG`, `BACKUP`, and `APPLICATION`. The value `ALL` specifies that root keys of all types are imported. If you do not specify any value for `<type>`, all key types are imported.

> **→ Recommendation**
>
> We recommend that you do not enter the password on the command line (`--password`). By omitting this parameter, you will be interactively prompted to enter it.

> **i Note**
>
> If you have backed-up root keys to different files, for example according to root key type, you need to execute the command several times.

- Using the SQL statement RECOVER ENCRYPTION ROOT KEYS
  The system database must be online to use this method.
  Execute the SQL statement in the system database:

```
RECOVER ENCRYPTION ROOT KEYS [ (<root_keytype_list>) ] FOR <database_name>
USING <root_key_backup_definition_file> PASSWORD <password>
```

  - `<database_name>` is the name of the tenant database.
  - The `<root_keytype_list>` option specifies the root key types and and accepts the values `PERSISTENCE`, `LOG`, `BACKUP`, and `APPLICATION`. If you do not specify any value for `<root_keytype_list>`, all root key types are imported.
  - `<root_key_backup_definition_file>` specifies the root key backup definition file in the file system.

Related Information

Stop a Tenant Database [page 384]
Recover a Database [page 628]

8.1.5 Auditing Database Activity

Auditing provides you with visibility on who did what in the SAP HANA database (or tried to do what) and when. This allows you, for example, to log and monitor read access to sensitive data.

Related Information

Activate and Configure Auditing [page 254]
Create an Audit Policy [page 255]
Delete Audit Entries [page 261]
Audit Trail Targets [page 269]
Best Practices and Recommendations for Creating Audit Policies [page 271]

8.1.5.1 Activate and Configure Auditing

The auditing feature of the SAP HANA database allows you to monitor and record selected actions performed in your database. To be able to use this feature, it must first be activated for the database. It is then possible to create and activate the required audit policies. You can do this on the *Auditing* page of the SAP HANA cockpit.

Prerequisites

- You have the system privileges `AUDIT ADMIN` and `INIFILE ADMIN`.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, click the *Auditing* card.

   The *Auditing* page opens.
2. Enable auditing.
   a. Open the *Configuration* tab and choose *Edit*.
   b. Set the auditing status to *Enable*.

   > **i** Note
   >
   > You can also enable auditing directly on *Auditing* card with the on/off switch.
3. Optional: Configure the required audit trail targets.

   By default, this is possible only in the **system database**.

   You can configure multiple audit trail targets: one for the database (*Overall Audit Trail Target*), and optionally one or more for the severity of audited actions, that is the audit level of the corresponding audit entries. If you do not configure a specific target for an audit level, audit entries are written to the overall audit trail target. For more information about the available audit trail targets, see *Audit Trail Targets*.

   Database table is the default audit trail target for tenant databases and *syslog* for the system database.

   > **i** Note
   >
   > If you are configuring auditing in a tenant database, you cannot change the audit trail targets. This is because the underlying system properties (`[auditing configuration] *_audit_trail_type`) are in the configuration change blacklist `multidb.ini`. Audit trails are by default written to an internal database table of the tenant database. Although not recommended, it is possible to change the audit trail target of a tenant database in the following ways:
   >
   > - The system administrator changes the audit trail targets for individual tenant databases directly by configuring the relevant system property (`[auditing configuration] *_audit_trail_type`) in the `global.ini` file. For more information about the system

properties for configuring audit trail targets and the configuration change blacklist in the *SAP HANA Security Guide*.

- The system administrator removes the relevant system property (`[auditing configuration] *_audit_trail_type`) from the configuration change blacklist, thus enabling the tenant database administrator to change the audit trail target.

> ⚠ Caution
>
> To ensure the privacy of tenant database audit trails, it is recommended that you do **not** change the default audit trail target (internal database table) of tenant databases.

4. Save your configuration.

Results

Auditing is activated in your database and you can now create audit policies.

Related Information

Create an Audit Policy [page 255]
Audit Trail Targets [page 269]
System Properties for Configuring Auditing
Default Blacklisted System Properties in Tenant Databases [page 393]

8.1.5.2 Create an Audit Policy

An audit policy defines the actions to be audited, as well as the conditions under which the action must be performed to be relevant for auditing. When an action occurs, the policy is triggered and an audit event is written to the audit trail. Audit policies are database specific. You can create audit policies on the *Auditing* page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege `AUDIT ADMIN`.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to the *Auditing* card, then click the card title.
2. On the *Audit Policies* tab of the *Auditing* page, click the *Create Audit Policy* button.
3. Enter the policy name.
4. Select the audited action status.

   The action status specifies when the actions in the policy are to be audited.

   The following values are possible:

Status	Description
All	All executions of the specified actions are audited, whether successful or unsuccessful.
Successful (default)	The action is audited only when the SQL statement is successfully executed.
Unsuccessful	The action is audited only when the SQL statement is unsuccessfully executed.

> **i Note**
>
> An unsuccessful attempt to execute an action means that the user was not authorized to execute the action. If another error occurs (for example, misspellings in user or object names and syntax errors), the action is generally not audited. In the case of actions that involve data manipulation (that is, INSERT, SELECT, UPDATE, DELETE, and EXECUTE statements), additional errors (for example, invalidated views) are audited.

5. Specify the actions to be audited.

   > **i Note**
   >
   > Only actions belonging to the same category can be combined together in the same policy.

   Selecting *All Actions* covers not only all actions that can be audited individually, but also actions that cannot otherwise be audited. Such a policy is referred to as a firefighter policy and is useful if you want to audit the actions of a particularly privileged user, for example.

   > **⚠ Caution**
   >
   > The actions that are audited are limited to those that take place inside the database engine while it is running. Therefore, database restart and recovery will not be audited.

6. If necessary, select 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.

7. If necessary, select 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 (*Only the selected users*). Alternatively, you can specify that the actions in the policy be audited when performed by all users **except** a particular user or users (*All users except those selected*).

   The actions in the policy will only be audited when performed by the specified user(s). If you do not specify a user, the actions will be audited regardless of who performs them.

   > **i** Note
   >
   > You **must** specify a user if you chose to audit all actions.

8. Optional: Select the audit level.

   The audit level specifies the severity of the audit entry written to the audit trail when the actions in the policy occur and ranges from INFO to EMERGENCY. The default level is INFO.

9. Optional: Select one or more policy-specific audit trail targets (system database only).

   Audit entries triggered by this policy will be written to the specified audit trail target(s). If you do not specify a policy-specific 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.

10. Optional: Specify a retention period.

    A retention period can only be specified for audit policies that have **database table** set **explicitly** as the audit trail target, not database table as the default target.

    Set *Delete audit entries automatically* to *On*.

    Specify a time period, after which audit entries will be deleted.

11. Specify whether you want the audit policy to be immediately enabled (default) or initially disabled on creation.
12. Choose *Review* to check the configuration of the new audit policy.
13. Save the new audit 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 polices 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 269]
System Properties for Configuring Auditing

8.1.5.2.1 Configure the Base Setup for Audit Policies

You can use a wizard to quickly apply a basic configuration for audit policies.

Prerequisites

- You have the system privilege `AUDIT ADMIN`.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Instead of manually configuring audit policies, you can use a wizard to apply SAP's recommended configuration settings. This allows you to quickly start working with audit policies.

For a registered database, the UI notifies you if audit policy configuration was not completed using the base setup wizards. You have the option to disable these notifications.

> **i** Note
>
> The base setup does not guarantee that the configuration is optimal for your particular system. To optimize the configuration, it may be advisable to manually configure specific settings.

Procedure

1. On the *Database Overview*, with the *Security and User Management* or *All* view selected, navigate to the *Auditing* card, then click the card title.
2. Choose *Base Setup*.

   Specify the base settings for auditing.

Option	Description
Display base setup notifications for Auditing	Enable or disable notifications for auditing. You can change this setting at any time.
Change Global Settings	To enable or disable all security-related notifications about the base setup, choose *User Interface Notification* > *Display base setup notifications*. You can enable or disable notifications at any time. For more information about SAP HANA cockpit user settings, see *Personalizing the Database Overview Page*.
Auditing Setup	To complete the base setup, choose *Complete Setup*. Follow the on-screen instructions. If you do not wish to change the configuration, choose *Manually set to completed* and choose *OK*. This disables notifications about the base setup.

You can configure the following settings for auditing:

- Auditing status
  It is recommended to enable auditing.
- Audit trail targets
  For more information, see *Audit Trail Targets*.
- Recommended audit policies
  You can choose which audit policies will be created.

3. Choose *Review* to display an overview of the auditing settings.

   To make changes, choose *Edit*.
4. Choose *Save*.

Results

The configuration settings are saved.

Notifications about the base setup are no longer displayed.

Related Information

Personalize the SAP HANA Cockpit [page 19]
Audit Trail Targets [page 269]
Create an Audit Policy [page 255]

8.1.5.3 Manage Audit Policies for Tenant Databases

You can change or delete audit policies for one or more tenant databases.

Prerequisites

- Auditing is enabled.
- You have the system privilege `DATABASE AUDIT ADMIN` in the system database.

Procedure

1. Connect to the system database.
2. On the *System Overview* page, choose *Manage Databases* > *Manage Audit Policies*.

   An overview of audit policies for all the tenant databases in the SAP HANA system appears.
3. Select an action.

Option	Description
To display the SQL statement for an audit policy...	Select one or more audit policies, then choose *Show SQL Statements*. If you select multiple audit policies, the SQL statements for each audit policy are displayed together.
To delete an audit policy...	Select one or more audit policies, then choose *Delete Audit Policies*. When you confirm, the policies are deleted and no more data is collected for them. To delete data that was collected by the audit policies, see *Delete Audit Entries*.

4. To change an audit policy, click its row.

   The configuration of the audit policy is displayed.
5. Choose *Edit*.
6. Make the required changes.

   For more information about configuration options for audit policies, see *Auditing Details*.
7. Select the tenant database(s) to be audited.

   In the *Audited Database* section, select one or more tenant databases in which to enable the policy.
8. Save.

Results

The changed policy appears in the list of audit policies. If you set the audit policy status to *Enabled*, the audit policy is automatically enabled in the tenant databases that you specified.

Related Information

Auditing Details [page 262]
Activate and Configure Auditing [page 254]
Create an Audit Policy [page 255]
Delete Audit Entries [page 261]

8.1.5.4 Delete Audit Entries

If the audit trail target is or was a database table, you can delete old audit entries, for example to prevent the audit table from growing indefinitely.

Prerequisites

- The audit trail target is or was *Database Table*.
- You have archived the audit entries that you plan to delete.
- You have the system privilege AUDIT OPERATOR.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

The database monitors the size of the table with respect to the memory allocation limit and issues an alert when it reaches defined values (by default 5%, 7%, 9%, and 11% of the allocation limit). This behavior can be configured with check 64.

> **i** Note
>
> If the table has grown so large that there is not enough memory available to delete old entries as described here, you can use the SQL command `ALTER SYSTEM CLEAR AUDIT LOG ALL` to completely empty the table. However, even if you archived the audit table beforehand (**recommended**), any new entries written between the time of archiving and the time of clearing may be lost.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, click the *Auditing* card.

   An overview of the audit policies is displayed.
2. Choose the *Audit Trail* tab.
3. Choose *Delete Audit Entries* and select which audit entries you want to be deleted.
   - Entries older than a specific number of days
   - Entries created before a date
   - All entries
4. Choose *Delete* to delete the specified audit entries from the audit table.

Results

The audit entries that you selected are deleted.

Related Information

Configure Alert Thresholds [page 131]

8.1.5.5 Auditing Details

On the *Auditing* page of the SAP HANA cockpit you can view and manage audit policies, audit trail targets, and the database table audit trail.

Audit Policies

> → Tip
>
> You can refine the list of audit polices by using the filtering options available in the table toolbar. Filter by policy name or audited action by entering the term directly in the search field, or click the *(Filter Settings)* button and select the required filter options. To clear all filters, click *(Clear All Filters)*.

Field	Description
Audit Policy	Audit policy name
Policy Status	An audit policy can be either *Enabled* or *Disabled*.
Audited Actions	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 *SAP HANA SQL and Systems View Reference*.
Audited Action Status	When the actions in the policy are to be audited: • On successful execution • On unsuccessful execution • All executions of the specified action are audited, whether successful or unsuccessful.
Audit Level	The severity of the audit entry written to the audit trail when the actions in the policy occur The following audit levels are possible • Emergency • Critical • Alert • Warning • Info
Users	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).
Audited Objects	The following audited object types are possible: • Schemas (and all objects contained within) • Tables • Views • Procedures • Sequences
Audit Trail Target	Policy-specific audit trail target(s) If there is no policy-specific audit trail target, audit entries generated by the policy are written to the audit trail target for the audit level of the policy if configured, or the audit trail target configured for the database. The applicable default audit trail target is always indicated in brackets.

Field	Description
Retention Period	Audit entries can be retained for a specified time, then deleted automatically. A retention period can only be specified for audit policies that explicitly have database table as the audit trail target. **i Note** If the retention period is changed, any audit entries that are no longer included are deleted immediately.

Configuration

Field	Description
Auditing Status	An audit policy can be either *Enabled* or *Disabled*.
Overall Audit Trail Target	The default audit trail target for the database 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.
Target for Audit Level Alert	The audit trail target to which audit entries with audit level `ALERT` are written
Target for Audit Level Emergency	The audit trail target to which audit entries with audit level `EMERGENCY` are written
Target for Audit Level Critical	The audit trail target to which audit entries with audit level `CRITICAL` are written

i Note

By default, it is not possible to configure audit trail targets in tenant databases. The audit trail target is *Database table*. For more information, see the section on audit trail targets.

Audit Trail

If the audit trail target is or was database table, you can view and manage the audit logs here, including the audit logs of the XS advanced run-time environment if available.

For more information, see the section on the audit trail view.

Related Information

Auditing Activity in the SAP HANA Database
Audit Trail Targets [page 269]
Audit Trail View [page 265]
Delete Audit Entries [page 261]

8.1.5.6 Audit Trail View

For each occurrence of an audited action, one or more audit entries are written to the audit trail. If the audit trail target is a database table, you can view the log on the *Auditing* page of the SAP HANA cockpit. It is also possible to view the audit logs of the XS advanced run-time environment.

To view the audit trail, on the *Database Overview* page, click the *Auditing* card.

On the *Audit Trail* tab, you can choose from the following log entries:

- *All Logs*
  Both database and XS advanced audit entries
- *SAP HANA Logs*
  Audit entries written by audit polices configured in the database
- *XSA Logs*
  Audit entries written by the audit log service of XS advanced

The following sections describe the layout of the audit trail and the options for configuring the layout.

Audit Trail Columns: SAP HANA Logs

Default Columns

Field	Description
Time Stamp	Time of event occurrence (in system local time)
Policy Name	Name of the audit policy that was triggered
Level	Severity of audited action
Status	Execution status of the statement
Client Host	Name of the host where the action occurred
User Name	User who performed the action
Statement	Statement that was executed

Additional Columns

Field	Description
Action	Action that was audited and thus triggered the policy

Field	Description
Application User Name	Application user who performed the action **⚠ Caution** Treat this information with caution. It comes from the application and SAP HANA has no way of verifying its authenticity.
Client Host	Name of the client machine
Client IP	IP address of the client application
Client PID	Process ID of the client process
Client Port	Port of the client process
Comment	Additional information about the audited event **i Note** Currently in case of failed logon attempts, the reason for failure appears in this field.
Connection ID	ID of the session in which the statement was executed
File Name	Configuration file name, for example global.ini
Grantable	Indication of whether the privilege or role was granted with or without GRANT/ ADMIN OPTION
Grantee	Name of the target user of the action, for example, grantee in a GRANT statement
Grantee Schema	Name of the schema of a granted or revoked role
Key	Configuration parameter, for example `global_auditing_state`
Object Name	Name of the object on which an action was performed, for example, a privilege was granted
Original Database	Name of tenant database in which the query originated; relevant for cross-database queries between tenant databases
Original User	Name of database user who executed the query in the origin tenant database; relevant for cross-database queries between tenant databases
Port	Port number
Previous Value	Previous value of the parameter, for example CSVTEXTFILE
Privilege Name	Name of the privilege that was granted or revoked
Role Name	Name of the role that was granted or revoked
Role Schema Name	Name of the schema in which a role was created/dropped, or the schema of a granted/ revoked role
Schema Name	Name of the schema where the action occurred, for example, a privilege was granted on a schema, or a statement was executed on object in a schema
Section	Configuration section name, for example auditing configuration
Service	Name of the service where the action occurred

Field	Description
Value	New parameter value, for example CSTABLE
XS Application User Name	XS application user name, for example, XSA_ADMIN

Audit Trail Columns: XSA Logs

Default Columns

Field	Description
Time	Time that the event occurred
XSA Organization ID	Application organization GUID
XSA Space ID	Application space GUID
App	Name of the application
XSA Instance ID	GUID of the used audit log service instance
User	Name of the application user
Message	Unique audit log message ID generated by the audit service

Additional Columns

Field	Description
Attachment Name	Name of the attachment that triggered the event.
Attachment ID	ID of the attachment that triggered the event
Attribute	Attribute that was changed
Attribute New Value	Old value of the attribute
Attribute Old Value	New value of the attribute
Binding ID	Application binding GUID in regards to the specific auditlog service instance that is being used
Category Type	
Client Host	IP of the client host
Client IP	IP of the client application
Client PID	PID of the client process
Client Port	Port of the client process
Connection ID	Connection ID
Created	Time of the event occurrence at the client side
Host	Name of the host where the event occurred
Level	Severity level of the event
Policy Name	Name of the audit policy triggered
Port	Port number

Field	Description
Service	Servcie name
Statement	SQL statement that caused the event
Statement User	Name of the user who executed the statement
Successful	Whether the event was successful or no
User Name	Name of the user connected to the database
XSA Channel	Communication protocol that was used when the audit event was triggered
XSA Data Subject	Owner of the accessed personal data
XSA Message IP	IP address of the event occurrence
XSA Object	Object containing the accessed personal data
XSA Tenant	XSA tenant GUID
XSA UUID	Unique audit log message ID generated by the audit service

Audit Trail View Configuration

You can configure the audit trail view by clicking the ⚙ *(Settings)* icon. The following options are available.

> **i Note**
>
> Your preferences are not saved when you log out of the cockpit.

Option	Description
Columns	Select the columns you want to see and the order in which they are displayed.
Sort	Sort the audit trail by one or more columns in ascending or descending order.
Filter	Filter the audit trail by creating complex include and exclude filters. **i Note** Values are case sensitive.
Group	Organize the audit trail by grouping events according to a particular field (for example, audited action, policy and so on).

Related Information

Delete Audit Entries [page 261]
Audit Trail Layout for Trail Target Database Table

8.1.5.7 Audit Trail Targets

In production systems, SAP HANA supports internal database table, syslog, and SAP HANA kernel trace as audit trail targets.

Audit Trail Target	Description
Internal database table	Using an SAP HANA database table as the target for the audit trail makes it possible to query and analyze auditing information quickly. It also provides a secure and tamper-proof storage location. Audit entries are only accessible through the public system views AUDIT_LOG, XSA_AUDIT_LOG, and the union of these two views ALL_AUDIT_LOG. Only SELECT operations can be performed on this view by users with the system privilege AUDIT OPERATOR, AUDIT ADMIN, or AUDIT READ. If a database table is explicitly configured as the audit trail target of an audit policy, you can define in the audit policy a retention period after which audit entries are automatically deleted. It is also possible to delete old audit entries by truncating the table. You can do this in SAP HANA cockpit or with the SQL statement ALTER SYSTEM CLEAR AUDIT LOG. The system monitors the size of the table with respect to the overall memory allocation limit of the system and issues an alert when it reaches defined values (by default 5%, 7%, 9%, and 11% of the allocation limit). This behavior can be configured with check 64 ("Total memory usage of table-based audit log"). Only users with the system privilege AUDIT OPERATOR can truncate the audit table.
Logging system of the Linux operating system (syslog)	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. **⚠ Caution** If the syslog daemon cannot write the audit trail to its destination, you will not be notified. To avoid a situation in which audited actions are occurring, but audit entries are not being written to the audit trail, ensure that the syslog is properly configured and that the audit trail target is accessible and has sufficient space available.

Audit Trail Target	Description
SAP HANA kernel trace	The audit log can be written to a kernel trace file (`*.ltc`) in the trace directory (`/usr/sap/<sid>/<instance>/<host>/trace`). The kernel trace output is not human-readable. It must be converted into a CSV files using the command-line tool `hdbltracediag` and then loaded into relational tables for SQL analysis. `hdbltracediag` is available on the SAP HANA server at `/usr/sap/<sid>/HDB<instance>/exe`.

Additionally, the option exists to store the audit trail in a CSV text file. This should only be used for test purposes in non-production systems. A separate CSV file is created for every service that executes SQL.

> ⚠ Caution
>
> You must not use a CSV text file for a production system as it has severe restrictions.
>
> Firstly, it is not sufficiently secure. By default, the file is written to the same directory as trace files (`/usr/sap/<sid>/<instance>/<host>/trace`). This means that database users with the system privilege DATA ADMIN, CATALOG READ, TRACE ADMIN, or INIFILE ADMIN can access it. In the SAP HANA database explorer, it is listed under *Database Diagnostics Files*, and at operating system level, any user in the SAPSYS group can access it.
>
> Secondly, audit trails are created for each server in a distributed database system. This makes it more difficult to trace audit events that were executed across multiple servers (distributed execution).

Audit Trails for Tenant Databases

By default, tenant database administrators **cannot** configure audit trail targets independently for their database since the underlying system properties are in the default configuration change blacklist (`multidb.ini`). The default target for all audit trails in tenant databases is internal database table. Although not recommended, it is possible to change the audit trail target of a tenant database in the following ways:

- The system administrator changes the audit trail targets for individual tenant databases directly by configuring the relevant system property (`[auditing configuration] *_audit_trail_type`) in the `global.ini` file. For more information about the system properties for configuring audit trail targets and the configuration change blacklist in the *SAP HANA Security Guide*.
- The system administrator removes the relevant system property (`[auditing configuration] *_audit_trail_type`) from the configuration change blacklist, thus enabling the tenant database administrator to change the audit trail target.

> ⚠ Caution
>
> To ensure the privacy of tenant database audit trails, it is recommended that you do **not** change the default audit trail target (internal database table) of tenant databases.

Related Information

System Properties for Configuring Auditing
Default Blacklisted System Properties in Tenant Databases [page 393]
Delete Audit Entries [page 261]
Configure Alert Thresholds [page 131]
ALTER SYSTEM CLEAR AUDIT LOG Statement (System Management)

8.1.5.8 Best Practices and Recommendations for Creating Audit Policies

General Best Practices

To reduce the performance impact of auditing, consider some basic guidelines for creating audit policies.

- Create as few audit policies as possible. It is usually better to have one complex policy than several simple ones.

  > → Remember
  >
  > Some audit actions cannot 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.
- Do not 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*.
- Do not 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;
```

- Changes to the SAP HANA license key

  Sample Code

  ```
  CREATE AUDIT POLICY "license creation" AUDITING ALL SET SYSTEM LICENSE
  LEVEL INFO;
  CREATE AUDIT POLICY "license deletion" AUDITING ALL UNSET SYSTEM LICENSE
  LEVEL INFO;
  ALTER AUDIT POLICY "license creation" ENABLE;
  ALTER AUDIT POLICY "license deletion" ENABLE;
  ```

- Recovery of tenant databases

  Sample Code

  ```
  CREATE AUDIT POLICY "recover database" AUDITING ALL RECOVER DATA LEVEL
  INFO;
  ALTER AUDIT POLICY "recover database" 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.

Related Information

Actions Audited by Default Audit Policy

CREATE AUDIT POLICY Statement (Access Control)

8.1.6 Configure the Database Password Policy and Password Blacklist

The passwords of database users are subject to certain rules. These are defined in the password policy and the password blacklist. You can change the default password policy of the database and maintain entries in the password blacklist in line with your organization's security requirements.

Prerequisites

- You have the system privilege INIFILE ADMIN.
- You have the object privileges SELECT, INSERT, and DELETE for the `_SYS_PASSWORD_BLACKLIST` table in the *_SYS_SECURITY* schema.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview*, with the *Security and User Management* or *All* view selected, click the *Authentication* card.

   The *Password Policy and Blacklist* page opens.
2. Click *Edit* in the footer bar.
3. In the *Password Policy* area, configure the options in line with your security requirements.

   All options have a default value. For more information about the individual options and their default values, see *Password Policy Details*.

   > → Tip
   >
   > To reset the default password policy, click *Set Default* in the footer toolbar. This will reset **all** options to their default values.
4. In the *Password Blacklist* area, add the words or partial words that you want to prohibit in passwords.

   The following configuration options are available:

Option	Description
Contained in Password	If you select this option, passwords that contain the blacklisted word are excluded. If you do not select this option, only passwords that match the blacklisted word exactly are excluded.
Case-Sensitive	If you select this option, the blacklisted word is case sensitive.

   > Example
   >
   > If you add the words **`SAP`**, **`my_sap_pwd`**, and **`sap_password`** to the blacklist and select the *Contained in Password* checkbox, then passwords containing "SAP", "my_sap_pwd", and "sap_password" are not allowed, regardless of how the password policy is configured.

5. Click *Save* to save the password policy and password blacklist.

Results

The passwords of database users must be created and changed in line with the defined policy.

Related Information

Configure a Password Policy for a User Group [page 283]

8.1.6.1 Password Policy Details

On the *Password Policy and Blacklist* page, you can view the password policy and change the its default configuration.

- Minimum Password Length [page 276]
- Lowercase Letters/Uppercase Letters/Numerical Digits/Special Characters Required [page 276]
- Lifetime of Initial Password [page 277]
- Minimum Password Lifetime [page 278]
- Maximum Password Lifetime [page 278]
- Maximum Duration of User Inactivity [page 279]
- Notification of Password Expiration [page 279]
- User Lock Time [page 279]
- Exempt SYSTEM User from Locking [page 280]
- Number of Allowed Failed Logon Attempts [page 280]
- Number of Last Used Passwords That Cannot Be Reused [page 281]
- Password Change Required on First Logon [page 282]
- Detailed Error Information on Failed Logon [page 282]

> **i Note**
>
> The individual password policy options are defined by parameters in the `password policy` section of the `indexserver.ini` configuration file, or in the case of the system database, the `nameserver.ini` file.

Minimum Password Length

The minimum number of characters that the password must contain

Parameter	`minimal_password_length`
Default Value	8 (characters)
Additional Information	You must enter a value between 6 and 64.
UI Label	*Minimum Password Length*

Lowercase Letters/Uppercase Letters/Numerical Digits/ Special Characters Required

The character types that the password must contain and how many

Parameter	`password_layout`
Default Value	Aa1 (at least one uppercase letter, at least one lowercase letter, and at least one number)

Additional Information	The following character types are possible: • Lowercase letter (a-z) • Uppercase letter (A-Z) • Numerical digits (0-9) • Special characters Any character that is not an uppercase letter, a lowercase letter, or a numerical digit is considered a special character. For example, underscore (_), hyphen (-). The following formats are supported for passwords: `<password> ::= { <letter> [ { <letter_or_digit> | # | $ }[…] ] | <digit> [ <letter_or_digit> […] ] | <any_quoted_string> }` If configuring this option in the `indexserver.ini` file using the `password_layout` parameter, you can use any specific letters, numbers and special characters, and the characters can be in any order. For example, the default value example could also be represented by **`a1A`**, **`hQ5`**, or **`9fG`**. To enforce the use of at least one of each character type including special characters, you specify **`A1a_`** or **`2Bg?`**. To enforce the use of a specific number of a particular character type, specify the character type multiple times. For example, if passwords must contain at least 3 digits, you could specify the layout with **`a123A`** or **`789fG`**. **i Note** Passwords containing special characters other than underscore must be enclosed in double quotes ("). When a password is enclosed in double quotes ("), any Unicode characters may be used. **⚠ Caution** The use of passwords enclosed in double quotes (") may cause logon issues depending on the client used. `hdsql` supports passwords enclosed in double quotes (").
UI Labels	*Lowercase Letters/Uppercase Letters/Numerical Digits/Special Characters Required*

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 **1**. 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. **i Note** In SAP HANA 1.0 SPS 12 and earlier, this parameter was misspelled as `maximum_unused_inital_password_lifetime`. If this parameter had a user-specified value before upgrade, this value will be set as the value of the parameter `maximum_unused_initial_password_lifetime`. The misspelled parameter is unset and disappears from the custom configuration file.
UI Label	*Lifetime of Initial Password*

Minimum Password Lifetime

The minimum number of days that must elapse before a user can change his or her password

Parameter	`minimum_password_lifetime`
Default Value	1 (day)
Additional Information	If you enter the value **0**, the password has no minimum lifetime.
UI Label	*Minimum Password Lifetime*

Maximum Password Lifetime

The number of days after which a user's password expires

Parameter	`maximum_password_lifetime`
Default Value	182 (days)
Additional Information	You must enter a value of at least **1**. A user administrator can exclude users from this password check with the following SQL statement: `ALTER USER <user_name> DISABLE PASSWORD LIFETIME`. However, this is recommended only for technical users only, not database users that correspond to real people. A user administrator can re-enable the password lifetime check for a user with the following SQL statement: `ALTER USER <user_name> ENABLE PASSWORD LIFETIME`.

UI Label	*Maximum Password Lifetime*

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 **1**. 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	*Maximum Duration of User Inactivity*

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	Notification is transmitted via the database client (ODBC or JDBC) and it is up to the client application to provide this information to the user. If you enter the value **0**, the user does not receive notification that his or her password is due to expire. The system also monitors when user passwords are due to expire and issues a medium priority alert (check 62). This may be useful for technical database users since password expiration results in the user being locked, which may affect application availability. It is recommended that you disable the password lifetime check of technical users so that their password never expires (`ALTER USER <technical_username> DISABLE PASSWORD LIFETIME`).
UI Label	*Notification of Password Expiration*

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	If you enter the value **0**, the user is unlocked immediately. This disables the functionality of parameter `maximum_invalid_connect_attempts`. A user administrator can reset the number of invalid logon attempts and reactivate the user account with the following SQL statement: `ALTER USER <user_name> RESET CONNECT ATTEMPTS`. To lock a user indefinitely, enter the value **-1**. On the *Password Policy and Blacklist* page of the SAP HANA cockpit, this corresponds to selecting the *Lock User Indefinitely* checkbox. The user remains locked until reactivated by a user administrator as described above.
UI Label	*User Lock Time*

Exempt SYSTEM User from Locking

Indicates whether or not the user SYSTEM is locked for the specified lock time (`password_lock_time`) after the maximum number of failed logon attempts (`maximum_invalid_connect_attempts`)

Parameter	`password_lock_for_system_user`
Default Value	true
Additional Information	This parameter cannot be configured for a user group.
UI Label	*Exempt SYSTEM User from Locking*

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	You must enter a value of at least **1**. A user administrator can reset the number of invalid logon attempts with the following SQL statement: `ALTER USER <user_name> RESET CONNECT ATTEMPTS` 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: • The number of invalid logon attempts since the last successful logon • The time of the last successful logon A user administrator can delete information about invalid logon attempts with the following SQL statement: `ALTER USER <user_name> DROP CONNECT ATTEMPTS` **→ Recommendation** 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. **i Note** Although this parameter is not valid for the SYSTEM user, the SYSTEM user will still be locked if the parameter `password_lock_for_system_user` is set to **`true`**. If `password_lock_for_system_user` is set to **`false`**, the SYSTEM user will not be locked regardless of the number of failed logon attempts.
UI Label	*Number of Allowed Failed Logon Attempts*

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	`last_used_passwords`
Default Value	5 (previous passwords)
Additional Information	If you enter the value **0**, the user can reuse his or her old password.
UI Label	*Number of Last Used Passwords That Cannot Be Reused*

Password Change Required on First Logon

Defines whether users have to change their initial passwords immediately the first time they log on

Parameter	`force_first_password_change`
Default Value	True
Additional Information	If this parameter is set to **`true`**, users can still log on with the initial password but every action they try to perform will return the error message that they must change their password. If this parameter is set to **`false`**, users are not forced to change their initial password immediately the first time they log on. However, if a user does not change the password before the number of days specified in the parameter `maximum_unused_initial_password_lifetime`, then the password still expires and must be reset by a user administrator. 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: `ALTER USER <user_name> FORCE PASSWORD CHANGE` A user administrator can override this password policy setting for individual users (for example, technical users) with the following SQL statement: • `CREATE USER <user_name> PASSWORD <password> [NO FORCE_FIRST_PASSWORD_CHANGE]` • `ALTER USER <user_name> PASSWORD <password> [NO FORCE_FIRST_PASSWORD_CHANGE]` **i Note** This parameter is only valid for users connecting with their SAP HANA database user name and password. It is not valid for connections established through other authentication mechanisms.
UI Label	*Password Change Required on First Logon*

Detailed Error Information on Failed Logon

Indicates the detail level of error information returned when a logon attempt fails

Parameter	`detailed_error_on_connect`
Default Value	false

Additional Information	If set to **`false`**, only the information `authentication failed` is returned. If set to **`true`**, the specific reason for failed logon is returned: • `Invalid user or password` • `User is locked` • `Connect try is outside validity period` • `User is deactivated`
UI Label	*Detailed Error Information on Failed Logon*

8.1.7 Configure a Password Policy for a User Group

If the users of a user group have different password requirements, you can configure group-specific values for the individual options of the password policy on the *User Groups* page of the SAP HANA cockpit.

Prerequisites

- You have the object privilege USERGROUP OPERATOR on the user group. If the user group is configured for shared administration (administration mode: *Group and user administrators*), system privilege USER ADMIN is also sufficient.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

The users of different user groups may have different requirements when it comes to passwords. For example, you may want the passwords of technical users to be very complex.

You configure group-specific values only for the required password policy options. You don't have to configure a value for all options. For those options without a group-specific value, the value from database password policy is simply copied. When you save and enable the password policy, it applies to users in the group until you enable the user group policy.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, choose the *User Group Management* link.
2. Select the relevant user group.
3. In the *Password Policy* area, choose *Configure* and configure the options in line with the user group's requirements.

The initial values for all options are copied from the password policy configured for the database. If you do not configure a group-specific value for a particular parameter, the value from database password policy applies. Choosing *Reset Default* reapplies all values from the database password policy.

For more information about the password policy options, see *Password Policy Details*.

4. To enable the group password policy, choose *Save and Enable*.

Results

The user group password policy is configured and effective for users in the user group.

> → Tip
>
> To determine which password policy a user is currently subject to, query the system view M_EFFECTIVE_PASSWORD_POLICY.

Related Information

User Groups
Password Policy Details [page 275]
Configure the Database Password Policy and Password Blacklist [page 274]
M_EFFECTIVE_PASSWORD_POLICY System View
Create a User Group [page 340]

8.1.8 Managing Client Certificates

SAP HANA uses X.509 client certificates as the basis for securing internal and external communication channels, as well as for several user authentication mechanisms. Certificates can be stored and managed in files in the file system and in some cases directly in the SAP HANA database.

Certificate Management in the Database

All certificate-based user authentication mechanisms in SAP HANA, as well as secure communication between SAP HANA and clients that access the SQL interface of the database rely on X.509 client certificates for authentication and verifying digital signatures. For ease of management, it's possible to store these certificates and configure their usage directly in the SAP HANA database.

In addition, in-database certificates must be used to secure communication during the process of copying or moving a tenant database between two systems, and to secure communication between SAP HANA and an LDAP server being used for user authentication and authorization.

The following figure shows for which purposes in-database certificates stored in certificate collections can be used. In-database certificates and certificate collections can be fully managed in the SAP HANA cockpit.

> **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.

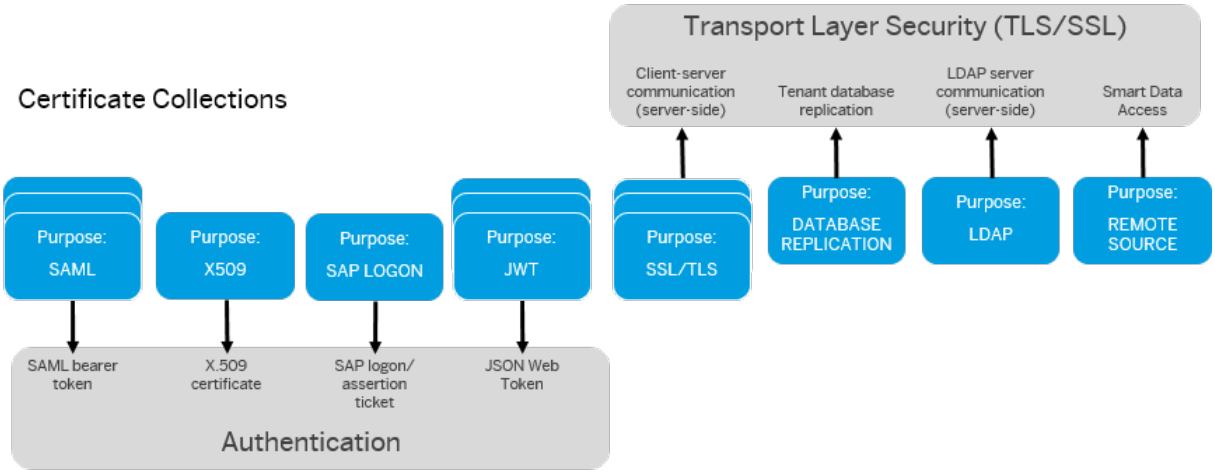

In-Database Certificates

Certificate Management in the File System

Although we recommend using in-database storage where possible, you can store and manage certificates in trust and key stores located in the file system, in so-called personal security environments or PSEs.

> **⚠ Caution**
>
> By default, the same PSE in the file system is shared by all databases for all external communication channels (including HTTP) and certificate-based authentication. Different PSEs must be explicitly configured for tenant databases.

> **→ Recommendation**
>
> You can migrate certificates from file-system based storage to in-database storage. If you do migrate certificates in the file system to the database, delete all related files from the file system to avoid any potential conflicts. For more information, see SAP Note 2175664.

However, not all certificates can be stored in the database, in particular the certificates required to secure internal communication channels using the system public key infrastructure (system PKI), and HTTP client access using SAP Web Dispatcher. These certificates are contained in PSE files located in the file system.

> **⚠ Caution**
>
> Do not delete these files from the file system.

The following figure shows for which purposes certificates stored in PSEs in the file system are possible. These PSEs are available by default and can be managed using for example the SAP Web Dispatcher administration

tool or the SAPGENPSE tool, both of which are delivered with SAP HANA. If you are using OpenSSL, you can also use the tools provided with OpenSSL.

> **i Note**
>
> OpenSSL is deprecated. If you are using OpenSSL, migrate to CommonCryptoLib. For more information, see SAP Note 2093286.

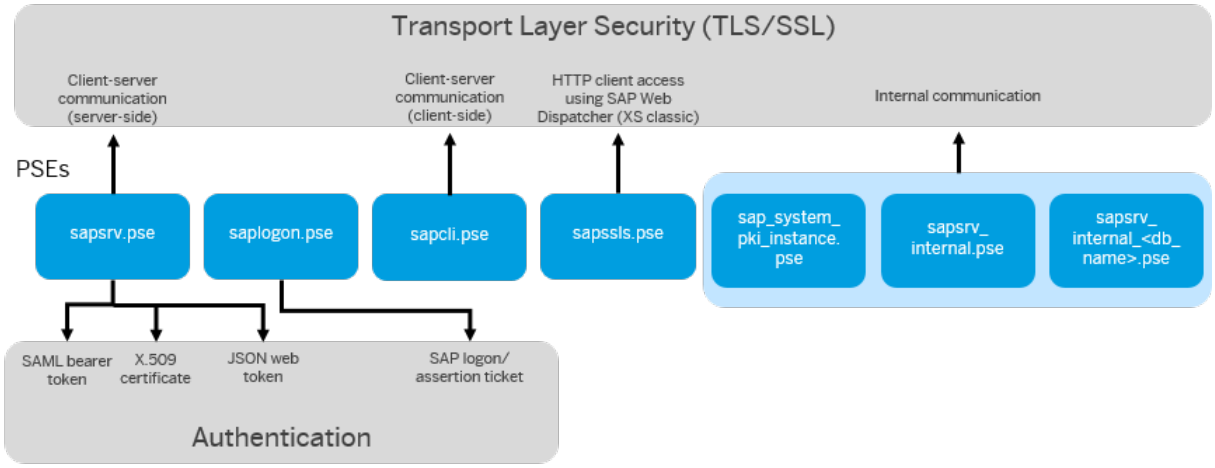

Default File-Based PSEs

Related Information

In-Database Certificate Management Workflow [page 287]
Client Certificates [page 288]
Certificate Collections [page 289]
SQL Statements and Authorization for In-Database Certificate Management (Reference) [page 301]
TLS/SSL Configuration on the SAP HANA Server
Server-Side TLS/SSL Configuration Properties for External Communication (JDBC/ODBC)
Copying and Moving Tenant Databases
SAP Note 2175664
SAP Note 2093286

8.1.8.1 In-Database Certificate Management Workflow

Managing certificates in the SAP HANA database follows a typical workflow. A full separation of duties is possible through user authorization.

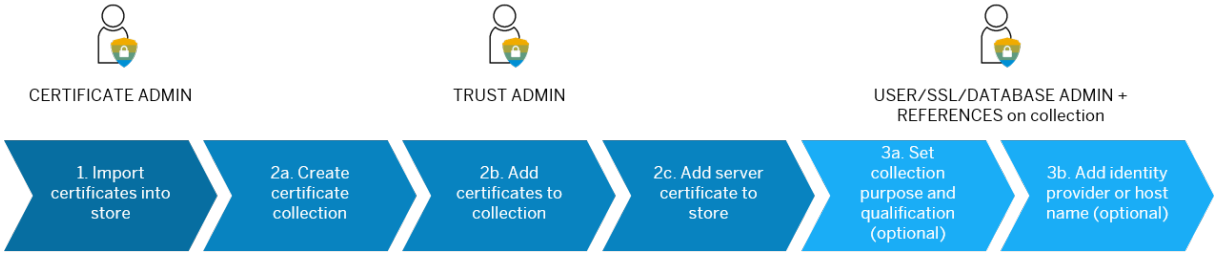

In-Database Certificate Management Workflow

1. A user with `CERTIFICATE ADMIN` system privilege imports into the certificate store the public-key certificates of trusted communication partners, as well as the root certificates of trusted Certification Authorities.
2. A user with `TRUST ADMIN` system privilege:
   1. Creates the required certificate collections.
   2. Adds trusted certificates from the certificate store to certificate collections.
   3. Adds the SAP HANA server certificate(s) to those collections that will be used for server authentication (for example, secure client-server communication over JDBC/ODBC).
3. A user with `USER ADMIN`, `SSL ADMIN`, or `DATABASE ADMIN` system privilege (depending on the purpose) and object privilege REFERENCES on the collection:
   1. Sets the purpose of individual collections with an optional qualification, for example, for one or more specific identity providers (purpose `JWT` or `SAML` only) or host names (purpose `SSL` only)
   2. Optional, at a later point in time: Adds one or more identity providers (purpose `JWT` or `SAML` only) or host names (purpose `SSL` only)

> **i Note**
>
> If you have qualified the purpose of a collection by adding an identity provider (purpose `SAML` or `JWT`) or a host name (purpose `SSL`), this collection takes precedence over a collection with the same purpose that has not been qualified.
>
> For example, if you have assigned your SAML provider SAML_PROVIDER to a certificate collection with the purpose SAML, only this collection is used to validate requests signed by SAML_PROVIDER. If you have other SAML providers defined in your system, requests from those providers are validated using the certificate collection with the purpose SAML that has not been qualified. If such a collection does not exist, then the corresponding PSE file in the file system is used for validation (by default, `sapsrv.pse`). For more information about changing the PSE file, see *Server-Side TLS/SSL Configuration Properties for External Communication (JDBC/ODBC)*.

Related Information

Import a Trusted Certificate into the Certificate Store [page 296]

Create a Certificate Collection [page 297]
Set the Purpose of a Certificate Collection [page 299]
Server-Side TLS/SSL Configuration Properties for External Communication (JDBC/ODBC)

8.1.8.2 Client Certificates

X.509 client certificates required for certificate-based user authentication and secure communication between SAP HANA and clients that access the SQL interface of the database can be stored and managed directly in the SAP HANA database.

Certificates stored in the SAP HANA database can be used for:

- Trust validation
  Certificates used for trust validation are the public-key certificates of trusted communication partners or root certificates from trusted Certification Authorities. These certificates contain the public part of a user's or component's public and private key pair.
- Server authentication
  Certificates used for server authentication are the public-key certificates of the SAP HANA server used to identify the server to connecting clients. In addition to the public-key information of the server, these certificates contain the server's private keys, as well as the intermediate certificates that complete the trust chain from the server certificate to the root certificate that the communication partner (client) trusts.

  > **i Note**
  >
  > Private keys are stored securely using the internal application encryption service of the SAP HANA database. For more information, see *Server-Side Data Encryption* in the *SAP HANA Security Guide*.

Once they have been imported into the database, certificates can be assigned to certificate collections. Certificate collections are also created and managed directly in the database, where they serve a unique purpose.

> **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 289]
Server-Side Data Encryption Services
CREATE CERTIFICATE Statement (System Management)
CERTIFICATES System View

8.1.8.3 Certificate Collections

A certificate collection (or PSE) is a secure location where the public information (public-key certificates) and private information (private keys) of the SAP HANA server are stored. A certificate collection may also contain the public information (public-key certificates) of trusted communication partners or root certificates from trusted Certification Authorities. Certificate collections can be created and managed as database objects directly in the SAP HANA database.

Certificate Collection Purposes

Certificate collections uniquely serve one of the following purposes in the database in which they exist.

For...	Set the Purpose...
Database replication with the aim of copying or moving a tenant database to another system	`DATABASE REPLICATION`
User authentication based on JSON Web Tokens	`JWT`
Communication between SAP HANA and an LDAP server being used for user authentication and authorization	`LDAP`
User authentication based on SAML assertions	`SAML`
User authentication based on logon and assertion tickets	`SAP LOGON`
Communication between SAP HANA and an LDAP server being used for user authentication and authorization	`SSL`
User authentication based on X.509 certificates	`X509`
Smart Data Access	`REMOTE SOURCE`

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.

> **i Note**
>
> You only need to set your own key or a private key in collections used for client-server communication, that is with purpose `SSL`. Although it is possible to do so, there is no need to set a private key for certificate collections used for trust validation.

> **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 with Qualified Purpose

Collections with the purposes `SAML`, `JWT`, and `SSL` may be qualified further. The purposes `SAML` and `JWT` can be specific to individual identity providers, and the purpose `SSL` can be specific to individual hosts.

This means that for SAML- and JWT-based authentication, you can assign the purpose `SAML` or `JWT` to several collections by adding different identity providers to each collection. In this way, the assertions or tokens issued by different providers can be validated separately. Identity providers must already exist in the database before they can be added to a collection, and each provider can only be assigned to one certificate collection.

Similarly, for client-server communication, you can assign the purpose `SSL` to several collections by adding host names to each collection. This allows you to configure separate sets of certificates for different hosts. Host names must only be entered in the correct format (for example, example.acme.com), and each host name can be assigned to several certificate collections.

If multiple collections with the same purpose exist, note the following behavior:

- Collections with qualified purposes take precedence over a collection with the same purpose that has not been qualified, that is a collection with the same purpose that has not been assigned any providers or hosts.

  > **i Note**
  >
  > Only one collection with an unqualified purpose may exist.

- If there is only one collection with a particular purpose and it is qualified, any unassigned providers or host names that may be configured cannot be validated. In this case, the trust store located on the file system and configured with the `global.ini` parameter `[communication] sslTrustStore` is used instead (`sapsrv.pse` by default).

This behavior can be illustrated with the following example.

> **Example**
>
> Multiple SAML identity providers are configured in your database (providerA, providerB, and providerC). Initially, two certificate collections are used for SAML-based user authentication. The first collection (pse1) is assigned the purpose `SAML` but no specific identity providers have been added. The second collection (pse2) is also assigned the purpose `SAML` but is qualified for a specific identity provider (providerA). This is configured as follows:
>
> ```
> SET PSE pse1 PURPOSE SAML;
> SET PSE pse2 PURPOSE SAML FOR PROVIDER providerA;
> ```
>
> With this configuration, collection pse2 has a qualified purpose and is used to validate requests signed by providerA only. Collection pse1 is not qualified with any specific providers and so is used to validate requests from all other providers configured in the database.
>
> If you now drop collection pse1, requests from providerB and providerC can no longer be validated. This is because there is no collection with the unqualified purpose `SAML`, and these providers have not been added to any other collection with purpose `SAML`. This means that it is no longer possible to log on or connect to the database using providerB and providerC (assuming that `sapsrv.pse` does not contain the necessary certificates).

In addition to creating a new collection with the unqualified purpose `SAML`, you could resolve the situation in a number of ways:

- Create one or more additional collections with the purpose `SAML` and assign the other providers to these:

  Sample Code

  ```
  CREATE PSE pse3;
  SET PSE pse3 PURPOSE SAML FOR PROVIDER providerB;
  CREATE PSE pse4;
  SET PSE pse4 PURPOSE SAML FOR PROVIDER providerC;
  ```

  Collections pse3 and pse4 are now used to validate requests signed by providerB and providerC respectively.
- Assign the other providers to the existing collection with qualified `SAML` purpose (pse2):

  Sample Code

  ```
  ALTER PSE pse2 ADD PROVIDER providerB, providerC;
  ```

  pse2 is now used to validate requests signed by providerA, providerB, and providerC.
- Remove the provider assigned to the existing collection, thus making it applicable to all providers:

  Sample Code

  ```
  SET PSE pse2 PURPOSE SAML;
  ```

  pse2 is now used to validate requests signed by all providers.

Ownership of Certificate Collections

A certificate collection is a database object created in runtime. It is therefore owned by the database user who creates it. If a certificate collection is in use, in other words it has been assigned one of the above purposes, it is not possible to change it (for example, add or remove certificates) or to delete it. However, if the owner of the certificate collection is deleted, the certificate collection will be deleted **even if it currently in use**.

⚠ Caution

The deletion of a certificate collection that is assigned a purpose could render the database unusable. For example, if TLS/SSL is being enforced for all client connections and the certificate collection used for TLS/SSL is deleted, no new client connections to the database can be opened.

Related Information

SAP HANA Authentication and Single Sign-On

Secure Communication Between SAP HANA and JDBC/ODBC Clients
Copying and Moving Tenant Databases
CREATE PSE Statement (System Management)
SET PSE Statement (System Management)

8.1.8.4 View Certificates in the Certificate Store

You can view certificates stored in the database on the *Certificate Store* page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege `CERTIFICATE ADMIN` or `TRUST ADMIN`.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to the *Security Related Links* card and choose the link *Certificate Store*.

The *Certificate Store* page opens. All certificates in the certificate store are listed. If you want to view the full details of a certificate, simply click it. For more information, see *Certificate Details*.

If the certificate is used in one or more certificate collections, you can navigate to the *Certificate Collections* page by clicking the collection name in the *Used In* column.

> **i Note**
>
> You will only see the certificate collection if you have the object privilege ALTER, DROP, or REFERENCES on the collection.

Related Information

Certificate Details [page 293]

8.1.8.4.1 Certificate Details

On the *Certificate Store* page of the SAP HANA cockpit you can view the details of all certificates in the certificate store of the SAP HANA database.

Field	Description
Issued To (CN)	Common name of the person or entity identified by the certificate
Issued To (DN)	Distinguished name of the person or entity identified by the certificate
Issued By (CN)	Common name of the entity that verified the information and issued the certificate
Issued By (DN)	Distinguished name of the entity that verified the information and issued the certificate
Issued On	Date on which the certificate was issued
Expires On	End of certificate's validity
Used In	The certificate collections to which the certificate has been assigned
Version	X.509 version (as specified in the corresponding RFC)
Public Key Algorithm	Public key algorithm
Public Key Length	Public key length
Signature Algorithm	The cryptographic algorithm used to sign the certificate
Basic Constraints	Whether the certificate belongs to a certification authority (CA)
Fingerprint	The hash of the entire certificate, used as a unique identifier in the certificate store
Serial Number	Serial number assigned by the certificate issuer

Related Information

Certificate Collection Details [page 295]

8.1.8.5 View Certificate Collections

You can view the certificate collections available in the database on the *Certificate Collections* page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege `CATALOG READ` and either `TRUST ADMIN`, `USER ADMIN`, or `SSL ADMIN`.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

On the *Database Overview* page, with the *Security and User Management* or *All* view selected, click the *Certificate Collections* link.
The *Certificate Collections* page opens. All existing collections are listed on the left. To see more detailed information about a specific collection on the right, simply select it. For more information, see *Certificate Collection Details*.

> **i Note**
>
> Certificate collections are also referred to as personal security environments (PSEs).

Related Information

Certificate Collections [page 289]

8.1.8.5.1 Certificate Collection Details

On the *Certificate Collections* page of the SAP HANA cockpit, you can view the details of all certificate collections in the SAP HANA database.

Field	Description
Purpose	Purpose of the collection • *DATABASE REPLICATION* Database replication for the purposes of copying or moving a tenant database to another system • *JWT* User authentication based on JSON Web Tokens • *LDAP* Communication between SAP HANA and an LDAP server being used for user authentication and authorization • *SAML* User authentication based on SAML assertions • *SAP LOGON* User authentication based on logon and assertion tickets • *SSL/TLS* Communication between SAP HANA and an LDAP server being used for user authentication and authorization • *X509* User authentication based on X.509 certificates
Provider	SAML or JWT identity provider if the collection purpose is *SAML* or *JWT*
Private Key	Indicates whether or not a private key has been set for the collection Only a collection with the purpose *SSL/TLS* requires a private key. This is the key that the SAP HANA server uses to identify itself to connecting clients. While possible, there is no need to set a private key for certificate collections used for trust validation.
Created By	Database user who created the collection
Comment	Optional comment

Field	Description
Certificates	Certificates assigned to the collection The function of each certificate in the certificate collection is indicated. The following functions are possible: • *TRUST* The certificate is the public-key certificate of a trusted communication partner. • *PERSONAL* The certificate is a server certificate belonging to the SAP HANA system and contains a private key. • *CHAIN* The certificate is an intermediate certificate that is part of the trust chain from the server certificate to the root certificate that the communication partner (client) trusts. For more information about the other certificate fields, see *Certificate Details*.

Related Information

Certificate Details [page 293]

8.1.8.6 Import a Trusted Certificate into the Certificate Store

You can store the public-key certificates of trusted communication partners, as well and the root certificates of trusted Certification Authorities directly in the SAP HANA database. You do this on the *Certificate Store* page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege System privilege `CERTIFICATE ADMIN`.
- The certificate that you want to add is available on your client in PEM format.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to the *Security Related Links* card and choose the link *Certificate Store*.
   The *Certificate Store* page opens, listing all certificates already in the certificate store.
2. Import the certificate:
   a. Click *Import*.
   b. Specify the location of the certificate file on your client or paste the content of the file.
   c. Click *OK*.

   The certificate is imported into the database and appears in the list of certificates in the certificate store. You can see the content of the certificate by navigating to its details view. For more information, see *Certificate Details*.

Results

The certificate is available for assignment to one or more certificate collections.

Related Information

Certificate Details [page 293]

8.1.8.7 Create a Certificate Collection

You can create a certificate collection on the *Certificate Collections* page. Then, you add the relevant trusted certificates and if necessary, the server certificate.

Prerequisites

- You have the system privilege `TRUST ADMIN`.
- The certificates you want to add to the collection are in the certificate store. For more information, see *Add a Certificate to the Certificate Store*.
- If you plan to add a server certificate to the collection, it is available on your client in PEM format.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to the *Security Related Links* card and choose the link *Certificate Collections*.

   The *Certificate Collections* page opens. All existing collections are listed on the left.

2. Create a new collection by clicking the + *(Add)* icon 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

Server-Side Data Encryption Services
Set the Purpose of a Certificate Collection [page 299]

8.1.8.8 Set the Purpose of a Certificate Collection

You specify the purpose of a collection on the *Certificate Collections* page, for example SAML user authentication. A collection may have only one purpose and a purpose may only be served by one collection.

Prerequisites

- If you are not the owner of the certificate collection, you need the object privilege `REFERENCES` on the certificate collection.
- You have the necessary system privilege to set the purpose:

Purpose	Privilege
User authentication	`USER ADMIN`
`SSL/TLS` (secure client-server communication over JDBC/ODBC)	`SSL ADMIN` **i Note** In addition, the server certificate containing the server's private key must be part of the collection.
`DATABASE REPLICATION`	`DATABASE ADMIN`
`LDAP`	`LDAP ADMIN`
`REMOTE SOURCE`	`CREATE REMOTE SOURCE`

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to the *Security Related Links* card and choose the link *Certificate Collections*.
2. Find and select the collection that you want to set the purpose for.
3. Open the collection for editing by choosing *Edit Purpose*.
4. In the dialog box, select the purpose:

Option	Description
DATABASE REPLICATION	Communication between two systems via external SQL connections for the purposes of copying or moving a tenant database
JWT	User authentication based on JSON Web Token (JWT)
LDAP	Communication between the SAP HANA database and an LDAP server being used for user authentication and authorization

Option	Description
SAML	User authentication based on SAML assertions
SAP LOGON	User authentication based on logon and assertion tickets
SSL/TLS	Client-server communication over JDBC/ODBC secured using SSL/TLS
X509	User authentication based on X.509 client certificates
REMOTE SOURCE	Access data stored in a remote data source. For example, using Smart Data Access or Rserve.

> **i** Note
>
> Only the purposes that have been enabled by the system administrator are visible, and only those you are authorized for are enabled.

5. Select a provider.

> **i** Note
>
> Only the providers that have been enabled by the system administrator are visible, and only those you are authorized for are enabled.

6. Specify a host.

   You need to manually specify the host name to be used.

7. 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.

8.1.8.9 Export a Client Certificate

You can export the contents of a client certificate available in the certificate store. For example, you may need to export the SAP HANA server certificate to set up a trust relationship with trusted clients.

Prerequisites

- You have the system privilege `CERTIFICATE ADMIN` or `TRUST ADMIN`.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

> **i** Note
>
> For information on how to set up the database client so that it accepts the server's certificate (or root certificate), including the use of the `openssl` or `sapgenpse` commands to extract and import these

certificates, refer to the section Implement Mutual Authentication in the *SAP HANA Client Interface Programming Reference*.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to the *Security Related Links* card and choose the link *Certificate Store*.

   The *Certificate Store* page opens. All certificates in the certificate store are listed.

   > **i Note**
   >
   > You can also navigate to certificates through the certificate collection to which they are assigned.

2. Find the certificate you want to export and navigate to the detailed view.
3. Click *Show PEM Representation* in the footer.
4. Export the certificate contents using copy and paste.

Related Information

Implement Mutual Authentication

8.1.8.10 SQL Statements and Authorization for In-Database Certificate Management (Reference)

All administration tasks related to in-database certificate management can be performed using SQL.

The following table lists the SQL statements for creating and managing certificates and certificate collections in the SAP HANA database, including the required authorization for each task. For more detailed information about the syntax of the statements mentioned, see the *SAP HANA SQL and System Views Reference*.

> **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	`SELECT * FROM CERTIFICATES` **i Note** You can also view certificates on the *Certificate Store* 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 certificate collections	`SELECT * FROM PSES` **i Note** You can also view certificate collections in the *Certificate Store* app of the SAP HANA cockpit.	System privilege TRUST ADMIN If you have object privilege ALTER, DROP, or REFERENCES on a certificate collection, you'll also be able to see this collection.
See which certificates are used in a certificate collection	`SELECT * FROM PSE_CERTIFICATES` **i Note** You can also see this information on the *Certificate Store* 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	`CREATE CERTIFICATE FROM <certificate_content> [ COMMENT <comment> ]`	System privilege CERTIFICATE ADMIN
Delete a certificate from the in-database certificate **i Note** If the certificate has already been added to a certificate collection, it can't be deleted.	`DROP CERTIFICATE <certificate_id>`	System privilege CERTIFICATE ADMIN
View certificate collections in the database, including the certificates they contain	`SELECT * FROM PSE_CERTIFICATES` **i Note** You can also view certificate collections on the *Certificate Collection* app of the SAP HANA cockpit.	System privilege CATALOG READ and either TRUST ADMIN, USER ADMIN, or SSL ADMIN **i Note** 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.
Create a certificate collection	`CREATE PSE <PSE_name>`	System privilege TRUST ADMIN

To...	Execute the Statement...	With the Authorization...
Add a public-key certificate to a certificate collection	`ALTER PSE <PSE_name> ADD CERTIFICATE <certificate_id>`	• Nothing if you're the owner of the certificate collection • Object privilege ALTER on the certificate collection if you're not the owner **i Note** If the purpose of the certificate collection has 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.
Remove a public-key certificate from a certificate collection	`ALTER PSE <PSE_name> DROP CERTIFICATE <certificate_id>`	
Add a private key to a certificate collection	`ALTER PSE <PSE_name> SET OWN CERTIFICATE <certificate_content>`	• Nothing if you're the owner of the certificate collection • Object privilege ALTER on the certificate collection if you're not the owner
Set the purpose of a certificate collection and qualify it if necessary	`SET PSE <PSE_name> PURPOSE <purpose>` The following PSE purposes are possible:	System privilege USER ADMIN if the purpose is user authentication (SAML, X.509, JWT, or logon tickets) **i Note** Object privilege REFERENCES on the certificate collection is additionally required if you are not the owner of the collection.

To...	Execute the Statement...	With the Authorization...
	SAML You can qualify the purpose SAML for specific identity providers as follows: `SET PSE <PSE_name> PURPOSE SAML FOR PROVIDER <provider, ...>` **i Note** A provider can only be used to qualify the purpose of one collection. If you use SET PSE to assign a provider that is already assigned to another collection, the provider is removed from the other collection. If this was the only provider assigned to the other collection, the purpose is also removed.	USER ADMIN
	JWT You can qualify the purpose JWT for specific identity providers as follows: `SET PSE <PSE_name> PURPOSE JWT FOR PROVIDER <provider, ...>` **i Note** A provider can only be used to qualify the purpose of one collection. If you use SET PSE to assign a provider that is already assigned to another collection, the provider is removed from the other collection. If this was the only provider assigned to the other collection, the purpose is also removed.	USER ADMIN
	X509	USER ADMIN
	SAP LOGON	USER ADMIN

To...	Execute the Statement...	With the Authorization...
	SSL You can qualify the purpose SSL for specific host names as follows: `SET PSE <PSE_name> PURPOSE SSL FOR HOST <host_name, ...>`. **i Note** The same host name can be used to qualify the purpose of multiple collections. **i Note** You can only assign the purpose SSL if a private key has already been added.	System privilege SSL ADMIN if the purpose is secure client-server communication (SSL)
	DATABASE REPLICATION	System privilege DATABASE ADMIN if the purpose is copying or moving a tenant database between systems
	LDAP	System privilege LDAP ADMIN if the purpose is LDAP-based user authentication and authorization
	REMOTE SOURCE	System privilege CREATE REMOTE SOURCE
Unset the purpose of a certificate collection	`UNSET PSE <PSE_name> PURPOSE <PSE_purpose>` **i Note** If the collection has identity providers or host names assigned, these are also removed from the collection.	

To...	Execute the Statement...	With the Authorization...
Add an identity provider or host name to a certificate collection with purpose SAML or JWT, or SSL.	• `ALTER PSE <PSE_name> ADD PROVIDER <SAML_or_JWT_provider>` • `ALTER PSE <PSE_name> ADD HOST <host_name>` **i Note** You can only use this statement to add further providers or host names to a collection. In other words, the collection must have at least one provider or host name already added. To add the first provider or host name, use the SET PSE statement.	• System privilege USER ADMIN if the purpose is SAML or JWT • System privilege SSL ADMIN if the purpose is secure client-server communication (SSL)
Remove an identity provider or host name from a certificate collection with purpose SAML or JWT, or SSL.	• `ALTER PSE <PSE_name> DROP PROVIDER <SAML_or_JWT_provider>` • `ALTER PSE <PSE_name> DROP HOST <host_name>` **i Note** You cannot use this statement to remove the last provider or host name assigned to a collection. To remove the last provider or host name, use the UNSET PSE statement.	• System privilege USER ADMIN if the purpose is SAML or JWT • System privilege SSL ADMIN if the purpose is secure client-server communication (SSL)
See which certificate collections have qualified purposes	`SELECT * FROM PSE_PURPOSE_OBJECTS`	System privilege TRUST ADMIN If you have object privilege ALTER or REFERNCES on a certificate collection or are the owner of the collection, you'll also be able to see information about that collection.
Delete a certificate collection	`DROP PSE <PSE_name>` **i Note** If the certificate collection has already been assigned a purpose, it can't be deleted.	• Nothing, if you're the owner of the certificate collection • Object privilege DROP on the certificate collection, if you're not the owner

Related Information

CREATE PSE Statement (System Management)
SET PSE Statement (System Management)
ALTER PSE Statement (System Management)
UNSET PSE Statement (System Management)
DROP PSE Statement (System Management)
CREATE CERTIFICATE Statement (System Management)
DROP CERTIFICATE Statement (System Management)
CERTIFICATES System View
PSES System View
PSE_CERTIFICATES System View

8.1.9 Data Anonymization

To enable analytics on data while still protecting the privacy of individuals, data anonymization capabilities are made available in SAP HANA in the form of anonymization views in the case of SQL, and as anonymization nodes in the case of calculation views.

> **i** Note
>
> SAP HANA provides only features and tools that help customers to implement data protection requirements and facilitate the required discussions between data scientists and data protection officers.

For more information about data anonymization in SAP HANA, see the *SAP HANA Security Guide*.

For more information about modeling calculation views with anonymization nodes, see the *SAP HANA Modeling Guide for SAP Web IDE for SAP HANA*.

Related Information

Show Anonymization Views [page 308]
Data Anonymization KPIs
KPI Parameters

8.1.9.1 Show Anonymization Views

As a data protection officer or data controller, you can retrieve information about all SQL anonymization views in the SAP HANA database as well as a list of all calculation views with anonymization nodes configured.

Prerequisites

- You have system privilege CATALOG READ.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to the *Anonymization Report* card and choose *View available anonymization views*.

The *Anonymization Report* page shows an overview of all the anonymized views in the database and all the calculation views for which one or more anonymization nodes is configured. The following information is shown:

- The name of the anonymized view or calculation view
- The name of the anonymization nodes (for calculation views only)
- The anonymization method used: *k*-anonymity, *l*-diversity, or differential privacy
- Configuration values of the relevant method
- Columns in the view, including anonymization information
- KPI values for anonymized views
  You can compare the KPI values for different versions of configuration for anonymization views.

Related Information

Data Anonymization KPIs
KPI Parameters

8.1.10 Display Information about an "Insufficient Privilege" Error

If an `Insufficient privilege` error occurs, you can find out more information about the missing privilege by using the associated GUID.

Prerequisites

To identify the missing privilege using a GUID, you need execute privilege for the following procedure:

```
SYS.GET_INSUFFICIENT_PRIVILEGE_ERROR_DETAILS('<GUID>', ?)
```

Context

If you encounter the error `Insufficient privilege`, you should inform the system administrator. The system administrator first needs to find out which privilege is needed by the user executing the statement. Then, the system administrator can decide whether the missing privilege should be assigned to the user.

Procedure

1. Make a note of the GUID shown in the error message.

   The following is an example of an error message:

   ```
   insufficient privilege: Detailed info for this error can be found with guid
   '3DFFF7D0CA291F4CA69B327067947BEE'
   ```
2. In SAP HANA cockpit, open the system with the missing privilege, and go to the *Database Overview*.
3. In either the *Security and User Management* or *All* view, navigate to the *Insufficient Privilege Details* card, specify the GUID and choose *Enter*.

   The missing privilege is displayed with the session user name and the checked user name. Optionally the object name, schema name, and object type are displayed.

   If the missing privilege is contained in one or more roles, the roles are displayed.

   > **i Note**
   >
   > If the missing privilege is an analytical privilege, neither the name of the privilege nor any roles can be displayed.

   For more information, see *Resolve Insufficient Privilege Errors* in the *SAP HANA Administration Guide (Security Administration and User Management)*.
4. Decide whether to assign the missing privilege or a role to the user.

Option	Description
To assign the missing privilege to the user...	Choose *Assign Privilege...*. You are prompted to select the privilege, and assign it using the privilege editor.
To assign a role containing the missing privilege...	Choose *Assign Role*. This function is only available if the current user is able to grant the role to the specified user.

Related Information

Configure Tracing in the SAP HANA Database Explorer
View Diagnostic Files in the SAP HANA Database Explorer
SAP HANA Administration Guide
Resolve Insufficient Privilege Errors
Resolve Insufficient Privilege Errors (Administration Guide)

8.2 User and Role Management

As a user administrator for the SAP HANA database, you create and configure database users, as well as authorize them to work with the database by assigning the necessary roles.

The recommended process for provisioning users is as follows:

1. Define and create roles.

   > **i Note**
   >
   > You can create only catalog roles using the SAP HANA cockpit. Design-time roles (database artifact with file suffix `.hdbrole`) can be created using the SAP Web IDE for SAP HANA and deployed using SAP HANA deployment infrastructure (SAP HANA DI).

2. Define and create user groups.
3. Create users in user groups.
4. Grant roles to users.

> **i Note**
>
> Creating user groups and assigning users to user groups is an optional step and depends on the requirements in your setup. For more information about user groups, see the *SAP HANA Security Guide*.

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
>
> Use SQL to create the technical user required to register a database through the SAP HANA cockpit and grant the minimum necessary authorizations:

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

Related Information

Create a Role [page 323]
Create a User Group [page 340]
Create a Database User [page 349]
Assign Roles to Database Users [page 358]

8.2.1 SAP HANA Database Users vs Cockpit Users

SAP HANA database users and cockpit users are two distinct users. Though they share similar concepts and terminology, such as roles and groups, they each have their own distinct meaning and function.

Each registered database has its own set of database users. A database user can monitor and manage a database in SAP HANA cockpit. Database users are managed in each registered database, using roles and privileges to control the tasks each can perform within the database. You can't log on to cockpit or Cockpit Manager using database credentials and database users do not appear within Cockpit Users in Cockpit Manager.

You use a cockpit user to log on to SAP HANA Cockpit or Cockpit Manager. Cockpit users are created and managed in Cockpit Manager. They are assigned cockpit roles, which dictate what tasks they can perform in Cockpit Manager and SAP HANA cockpit. Cockpit users are assigned to database groups, which contain the registered databases that will be available in SAP HANA cockpit. Once logged in to cockpit, you can't access one of the registered databases unless you supply the credentials of a database user. Cockpit users do not appear in User Management in a registered database in cockpit.

8.2.2 View a Database User

You can view database users on the *User Management* page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege CATALOG READ. You don't require any additional privileges to view your own database user.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *User Management*.

   An overview of all existing database users is displayed.
2. To see more detailed information about a specific user, simply select it.

   For more information, see *Database User Details*.

Related Information

Database User Details [page 312]

8.2.2.1 Database User Details

On the *User Management* page of the SAP HANA cockpit, you can view the details of all users in the SAP HANA database.

General Information

Field	Description
User Name	Unique user name

Field	Description
User Group	A user can belong to only one user group, but a user does not have to be a member of any group.
E-Mail	User's e-mail address
Valid From/To	Validity period of the user If the user account is not currently within its validity period, the user is inactive and cannot log on. If no validity period is configured, the user is indefinitely valid.
Last Successful Login	The last time the user logged on successfully
Deactivated Since	Time the user was deactivated
Creation of Objects in Own Schema	Indicates whether or not the user can create objects in their database schema Standard users can create objects in their schema. Restricted users cannot. For more information about the difference between standard and restricted users, see *Database Users*.
PUBLIC Role	Indicates whether or not the user has the PUBLIC role Standard users have this role by default. Restricted users do not. For more information about the difference between standard and restricted users, see *Database Users*.
Disable ODBC/JDBC Access	Indicates whether or not the user can connect to the database via ODBC or JDBC By default, ODBC/JDBC access is disabled for restricted users, meaning they can only connect via HTTP/HTTPS, and enabled for standard users. For more information about the difference between standard and restricted users, see *Database Users*.
Comment	Free-text comment or description (if applicable)

Authorization Mode

Field	Description
Authorization Mode	Indicates whether the user's authorization is based on LDAP group membership or local SAP HANA mechanisms A user with authorization mode *LDAP* is granted roles exclusively based on their LDAP group membership. It is not possible to grant such a user other roles or privileges directly. The default user authorization mode is *Local*. This means that the user must be granted roles and privileges directly. For more information about LDAP group authorization, see the *SAP HANA Security Guide*.
Assign Roles	Open the *Assign Roles* app where you can assign roles to the user
Assign Privileges	Open the *Assign Privileges* app where you can assign privileges to the user

Authentication

Field	Description
Password	Indicates whether or not user name-password authentication is enabled Users accessing the SAP HANA database authenticate themselves by entering their database user name and either their local SAP HANA password or a password stored in an LDAP directory server.
Force password change on next logon	Indicates whether the user must change a password set by a user administrator the first time he or she logs on, regardless of how the password policy parameter *Password Change Required on First Logon* is configured
Kerberos	Indicates whether or not Kerberos authentication is enabled If enabled, the external identity to which the database user is mapped must be specified.
SAP Logon Ticket, *SAP Assertion Ticket*	Indicates whether or not authentication using SAP logon or assertion tickets is enabled
SAML	Indicates whether or not SAML authentication is enabled If enabled, the external identity to which the database user is mapped can be explicitly specified. Alternatively, if the option *Automatic Mapping by Provider* is selected, the identity provider is allowed to map its users to the database user.

Field	Description
JWT	Indicates whether or not JSON Web Token authentication is enabled If enabled, the external identity to which the database user is mapped can be explicitly specified. Alternatively, if the option *Automatic Mapping by Provider* is selected, the identity provider is allowed to map its users to the database user.
X509	Indicates whether or not X.509 certificate authentication is enabled

> **i** Note
>
> For more information about authentication mechanisms, see the *SAP HANA Security Guide*.

Custom User Properties

Additional user properties can be configured for client applications. The following properties are available by default:

Property	Description
CLIENT	The session client When you create SAP HANA calculation views, it is possible to filter the data according to the client specified in table fields such as MANDT or CLIENT.
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 Priority values of 0 (lowest priority) to 9 (highest) are available; the default priority is 5.

Property	Description
STATEMENT MEMORY LIMIT	The maximum memory (in GB) that can be used by a statement executed by the user The properties `statement_memory_limit` and `statement_memory_limit_threshold` in the `memory_manager` section of the `global.ini` configuration file are used to limit the memory that can be allocated with respect to statement execution. `statement_memory_limit_threshold` indicates what percentage of the global memory allocation limit must be in use before the specific value of `statement_memory_limit` 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 *Setting a Memory Limit for SQL Statements* in the *SAP HANA Administration Guide*.
STATEMENT THREAD LIMIT	The maximum number of threads that can be used by a statement executed by the user
TIME ZONE	The user's timezone The standard database formats for locale and timezone are supported.

Related Information

User Groups
Create a User Group [page 340]
Filter Data for Specific Clients
Setting a Memory Limit for SQL Statements
LDAP Group Authorization
Configure LDAP Group Authorization
Assign Roles to Database Users [page 358]
Assign Privileges to a User [page 360]
User Authentication Mechanisms
SAP HANA Security Guide

8.2.2.2 Database Users

Every user who wants to work with the SAP HANA database must have a database user.

Database users are created with either the `CREATE USER` or `CREATE RESTRICTED USER` statement, or using SAP HANA cockpit (*User Managerment*).

Standard Users

Standard users correspond to users created with the `CREATE USER` statement. By default they can create objects in their own schema and read data in system views. Read access to system views is granted by the `PUBLIC` role, which is granted to every standard user.

Restricted Users

Restricted users, which are created with the `CREATE RESTRICTED USER` statement, initially have no privileges. Restricted users are intended for provisioning users who access SAP HANA through client applications and who are not intended to have full SQL access via an SQL console. If the privileges required to use the application are encapsulated within an application-specific role, then it is necessary to grant the user only this role. In this way, it can be ensured that users have only those privileges that are essential to their work.

Compared to standard database users, restricted users are initially limited in the following ways:

- They cannot create objects in the database as they are not authorized to create objects in their own database schema.
- They cannot view any data in the database as they are not granted the standard PUBLIC role.
- They are only able to connect to the database using HTTP/HTTPS.
  For restricted users to connect via ODBC or JDBC, access for client connections must be enabled by executing the SQL statement `ALTER USER <user_name> ENABLE CLIENT CONNECT` or enabling the corresponding option for the user in the SAP HANA cockpit.
  For full access to ODBC or JDBC functionality, users also require the predefined role `RESTRICTED_USER_ODBC_ACCESS` or `RESTRICTED_USER_JDBC_ACCESS`.

  > **i Note**
  >
  > Disabling ODBC/JDBC access for a user, either a restricted user or a standard user, does not affect the user's authorizations or prevent the user from executing SQL commands via channels other than JDBC/ODBC. If the user has been granted SQL privileges (for example, system privileges and object privileges), he or she is still authorized to perform the corresponding database operations using, for example, a HTTP/HTTPS client.

A user administrator can convert a restricted user into a standard user (or vice versa) as follows:

- Granting (or revoking) the PUBLIC role
  You can do this by editing the user in the SAP HANA cockpit or with the SQL statement `ALTER USER <username> GRANT | REVOKE ROLE PUBLIC`.
- Granting (or revoking) authorization to create objects in the user's own schema
  You can do this by editing the user in the SAP HANA cockpit or with the SQL statement `ALTER USER <username> GRANT | REVOKE CREATE ANY ON OWN SCHEMA`.
- Enabling (or disabling) full SQL
  You can do this by editing the user in the SAP HANA cockpit or with the SQL statement `ALTER USER <user_name> ENABLE CLIENT CONNECT`.

> **Note**
>
> A user is only identified as a restricted user in system view USERS if he doesn't have the PUBLIC role or authorization for his own schema.

Predefined Database Users

When an SAP HANA database is created, several database users are created by default. The most important of these is the `SYSTEM` database user, which should be deactivated in production systems.

Several technical database users (that is, database users that do not correspond to real people) are also created, for example, `SYS` and `_SYS_REPO`.

For more information about other predefined database users, see the *SAP HANA Security Guide*.

Related Information

Predefined Users

8.2.3 View a Database Role

You can view database roles on the *Role* page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege CATALOG READ or ROLE ADMIN.

  > **Note**
  >
  > Even if ROLE ADMIN was revoked from your user, you can still view roles that you created yourself.

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *Role Management*.

   An overview of all existing database roles is displayed.

2. To see more detailed information about a specific role, simply select it.

   For more information, see *Database Role Details*.

Related Information

Database Role Details [page 319]

8.2.3.1 Database Role Details

On the *Role Management* page of SAP HANA cockpit, you can view the details of all roles in the SAP HANA database.

General Information

Field	Description
Schema	Schema in which the role exists (if applicable) The schema represents the role's runtime namespace, which allows it to be used in different contexts. A role without a schema is a global role. **⚠ Caution** A role with a namespace will be deleted if the schema is deleted.
Creator	User who created the role
LDAP Groups	LDAP groups that have been mapped to the role (if applicable) Database users configured for LDAP authorization who belong to the specified group(s) are automatically granted the role in line with your LDAP configuration for SAP HANA. For more information, see the *SAP HANA Security Guide*.
Comment	Free-text comment or description (if applicable)

Field	Description
Type	The role type: • *Catalog* A role created in run-time with the SQL statement `CREATE ROLE` • *Catalog (LDAP)* A catalog role with LDAP group mappings • *HDI* A role created using the SAP Web IDE for SAP HANA and deployed using SAP HANA deployment infrastructure (SAP HANA DI) • *HDI (LDAP)* A HDI role with LDAP group mappings LDAP groups are mapped to the activated catalog role. • *Repository role* A role created in the built-in repository of the SAP HANA database using either the SAP HANA Web Workbench or the SAP HANA studio
Is Part of Roles	Indicates whether the role is included in another role
Roles	Other roles granted to this role
System Privileges	System privileges granted to the role 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. For a list of all system privileges, see *System Privileges (Reference)* in the *SAP HANA Administration Guide*.
Object Privileges	Object privileges granted to the role 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). For a list of all object privileges, see *Object Privileges (Reference)* in the *SAP HANA Administration Guide*.
Analytic Privileges	Analytic privileges granted to the role Analytic privileges are used to control read access to data in SAP HANA information models (that is, analytic views, attribute views, and calculation views) depending on certain values or combinations of values.
Is Part of Roles	Other roles to which this role has been granted

Field	Description
Application Privileges	Application privileges granted to the role Application privileges are used to authorize user and client access to SAP HANA XS classic applications. **i Note** Application privileges are not relevant in the context of SAP HANA XS advanced applications. For more information about the authorization concept of the SAP HANA XS advanced, see the *SAP HANA Security Guide*.
Package Privileges	Package privileges granted to the role Package privileges are used to allow access to and the ability to work in packages in the repository of the SAP HANA database. **i Note** With SAP HANA XS advanced, source code and web content are not versioned and stored in the repository, so package privileges are not relevant in this context.
Privileges on Users	Privileges on users granted to the role 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.

Related Information

Database Roles [page 322]
System Privileges
System Privileges (Reference)
Object Privileges
Object Privileges (Reference)
Analytic Privileges
Package Privileges
Application Privileges
LDAP Group Authorization
SAP HANA Security Guide

8.2.3.2 Database Roles

A database role is a collection of privileges that can be granted to either a database user or another role in runtime.

A role typically contains the privileges required for a particular function or task, for example:

- Business end users reading reports using client tools such as Microsoft Excel
- Modelers creating models and reports
- Database administrators operating and maintaining the database and its users

Privileges can be granted directly to users of the SAP HANA database. However, roles are the standard mechanism of granting privileges as they allow you to implement complex, reusable authorization concepts that can be modeled on business roles.

Creation of Roles

Roles in the SAP HANA database can exist as runtime objects only (catalog roles), or as design-time objects that become catalog objects on deployment (database artifact with file suffix `.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, or HDI).

> **i** Note
>
> Due to the container-based model of HDI where each container corresponds to a database schema, HDI roles, once deployed, are schema specific.

SAP HANA XS advanced has the additional concept of application roles and role collections. These are independent of database roles in SAP HANA itself. In the XS advanced context, SAP HANA database roles are used only to control access to database objects (for example, tables, views, and procedures) for XS advanced applications. For more information about the authorization concept of XS advanced, see the *SAP HANA Security Guide*.

Role Structure

A role can contain any number of the following privileges:

- **System privileges** for general system authorization, in particular administration activities
- **Object privileges** (for example, SELECT, INSERT, UPDATE) on database objects (for example, schemas, tables, views, procedures, and sequences)
- **Analytic privileges** on SAP HANA information models
- **Package privileges** on repository packages (for example, `REPO.READ`, `REPO.EDIT_NATIVE_OBJECTS`, `REPO.ACTIVATE_NATIVE_OBJECTS`)

- **Application privileges** for enabling access to SAP HANA-based applications developed in an SAP HANA XS classic environment

> **i** Note
>
> There are no HDI or XS advanced equivalents in the SAP HANA authorization concept for package privileges on repository packages and applications privileges on SAP HANA XS classic applications. For more information about the authorization concept of XS advanced, see the *SAP HANA Security Guide*.

A role can also contain other roles.

Roles Best Practices

For best performance of role operations, in particular, granting and revoking, keep the following basic rules in mind:

- Create roles with the smallest possible set of privileges for the smallest possible group of users who can share a role (principle of least privilege).
- Avoid granting object privileges at the schema level to a role if only a few objects in the schema are relevant for intended users.
- Avoid creating and maintaining all roles as a single user. Use several role administrator users instead.

Related Information

Authorization in SAP HANA XS Advanced

8.2.4 Create a Role

You can create a new role and grant it privileges and roles on the *Role Management* page of the SAP HANA cockpit. You can also map roles to LDAP groups if you are implementing user authorization based on LDAP group membership.

Prerequisites

- You have the system privilege ROLE ADMIN.
- You have the privileges required to grant privileges and roles to the new role.
  For more information, see *Prerequisites for Granting and Revoking Privileges and Roles*.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *Role Management*.

   An overview of all existing database roles is displayed.
2. Create the role:
   a. Click the + *(Add)* icon.
   b. Specify a unique role name.
   c. Optional: Enter a comment or text to describe the role.
   d. Optional: Assign the role a runtime namespace by choosing the schema in which to create the role.

      Role namespaces allow you to reuse role names in different contexts. If you do not select a schema, the role will be created as a global role.

      > ⚠ Caution
      >
      > A role with a namespace will be deleted if the schema is deleted.
3. Optional: If you are implementing user authorization based on LDAP group membership, map one or more LDAP group to the role:
   a. Enable the assignment of LDAP groups to the role.
   b. Add the required LDAP groups by specifying the unique distinguished name (DN).

   Users configured for LDAP authorization who belong to the specified group(s) are automatically granted the role in line with your LDAP configuration for SAP HANA. For more information, see the *SAP HANA Security Guide*.
4. Choose *Create*.

   The role is created.
5. Assign the required roles to the role:
   a. In the *Roles* area, choose *Edit*.
   b. Choose *Add* and select the roles you want to assign.
   c. If you want users who have the new role to be able to grant the assigned role on to others, choose *Grantable to Others*.
   d. Save the role assignment.
6. Assign the required privileges to the role:
   a. For the relevant privilege type, choose *Edit*.
   b. Choose *Add* and select the privileges you want to assign.

      > ℹ Note
      >
      > For object and package privileges, you must first add the object or package and then add the required privilege to the object or package.
   c. If you want users who have the new role to be able to grant the assigned privilege on to others, choose *Grantable to Others*.
   d. Save the privilege assignment.
   e. Repeat for further privilege types.

Results

The role is created and appears in the list of roles on the left.

Next Steps

Assign the role to the required database users. You should only do this for users with authorization mode *Local*, not *LDAP*.

If you mapped LDAP groups to the role, configure the connection to the LDAP provider and configure the required database users for LDAP group authorization. For more information, see the *SAP HANA Administration Guide*.

Related Information

Prerequisites for Granting and Revoking Privileges and Roles [page 325]
Privileges
System Privileges (Reference) [page 328]
Object Privileges (Reference) [page 333]
Database Role Details [page 319]
Assign Roles to Database Users [page 358]
Configure LDAP Group Authorization
LDAP Group Authorization
SAP HANA Security Guide
SAP HANA Administration Guide

8.2.4.1 Prerequisites for Granting and Revoking Privileges and Roles

To be able to grant and revoke privileges and roles to and from users and roles, several prerequisites must be met.

The following table lists the prerequisites that a user must meet to grant privileges and roles to another user (or role).

Prerequisites for Granting Privileges

To grant...	The granting user needs...
A system privilege	The system/object privilege being granted and be authorized to grant it to other users and roles

To grant...	The granting user needs...
An object privilege on an object that exists only in runtime	
An object privilege on an activated object created in the repository, such as a calculation view	The object privilege EXECUTE on the procedure GRANT_PRIVILEGE_ON_ACTIVATED_CONTENT
An object privilege on schema containing activated objects created in the repository, such as a calculation view	The object privilege EXECUTE on the procedure GRANT_SCHEMA_PRIVILEGE_ON_ACTIVATED_CONTENT
A package privilege	The package privilege being granted and be authorized to grant it to other users and roles
An analytic privilege	The object privilege EXECUTE on the procedure GRANT_ACTIVATED_ANALYTICAL_PRIVILEGE
An application privilege	The object privilege EXECUTE on the procedure GRANT_APPLICATION_PRIVILEGE
Privilege on user ATTACH DEBUGGER	To be the user on which ATTACH DEBUGGER is granted
A role created in runtime	Either: • The role being granted and be authorized to grant it to other users and roles, or • The system privilege ROLE ADMIN
A role created in the repository	The object privilege EXECUTE on the procedure GRANT_ACTIVATED_ROLE
A role created in an HDI container	Either: • Privileges to execute GRANT_CONTAINER_SCHEMA_ROLES in the container's API schema, or, if the user is a container group administrator, privileges to execute GRANT_CONTAINER_SCHEMA_ROLES in the container group's API schema. • The system privilege ROLE ADMIN

Prerequisites for Revoking Privileges

To revoke ...	The revoking user needs...
A system privilege	To be the user who granted the privilege
An object privilege on an object that exists only in runtime	
An object privilege on an activated object created in the repository, such as a calculation view	The object privilege EXECUTE on the procedure REVOKE_PRIVILEGE_ON_ACTIVATED_CONTENT
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

To revoke ...	The revoking user needs...
Privilege on user ATTACH DEBUGGER	To be the user on which ATTACH DEBUGGER is granted
A role created in runtime	• To be the user who granted the role, or • The system privilege ROLE ADMIN **i Note** With the exception of roles granted by technical user _SYS_REPO, a user with ROLE ADMIN **cannot** revoke roles granted by technical users SYS and _SYS*.
A role created in the repository	The object privilege EXECUTE on the procedure REVOKE_ACTIVATED_ROLE
A role created in an HDI container	Either: • Privileges to execute REVOKE_CONTAINER_SCHEMA_ROLES in the container's API schema, or, if the user is a container group administrator, privileges to execute REVOKE_CONTAINER_SCHEMA_ROLES in the container group's API schema. • The system privilege ROLE ADMIN

Authorization of User _SYS_REPO

If you are implementing your authorization concept using roles and you are creating roles in the repository of the SAP HANA database, the technical user _SYS_REPO is the granting and revoking user. _SYS_REPO **automatically** meets all of the above prerequisites with the exception of those for granting/revoking objects privileges on objects that exist only in runtime. These privileges must be explicitly granted to _SYS_REPO. For more information about roles as repository objects, see the *SAP HANA Security Guide*.

How are HDI roles granted and revoked?

In the SAP HANA deployment infrastructure (HDI), there are two types of users that can grant or revoke roles.

Roles created in an HDI container can be granted to (or revoked from) other roles and users either by an administrator of the container or an administrator of the container group that the container belongs to.

Preferably, roles created in an HDI container are granted or revoked by a container administrator.

Related Information

Create and Authorize a User
Create and Authorize a Restricted User
Repository Roles
SAP HANA Security Guide

8.2.4.2 System Privileges (Reference)

System privileges control general system activities.

General System Privileges

System Privilege	Description
ADAPTER ADMIN	Controls the execution of the following adapter-related statements: CREATE ADAPTER, DROP ADAPTER, and ALTER ADAPTER. It also allows access to the ADAPTERS and ADAPTER_LOCATIONS system views.
AGENT ADMIN	Controls the execution of the following agent-related statements: CREATE AGENT, DROP AGENT, and ALTER AGENT. It also allows access to the AGENTS and ADAPTER_LOCATIONS system views.
ATTACH DEBUGGER	Authorizes debugging across different user sessions. For example, userA can grant ATTACH DEBUGGER to userB to allow userB to debug a procedure in userA's session (userB still needs DEBUG privilege on the procedure, however).
AUDIT ADMIN	Controls the execution of the following auditing-related statements: CREATE AUDIT POLICY, DROP AUDIT POLICY, and ALTER AUDIT POLICY, as well as changes to the auditing configuration. It also allows access to the AUDIT_LOG, XSA_AUDIT_LOG, and ALL_AUDIT_LOG system views.
AUDIT OPERATOR	Authorizes the execution of the following statement: ALTER SYSTEM CLEAR AUDIT LOG. It also allows access to the AUDIT_LOG system view.
AUDIT READ	Authorizes read-only access to the rows of the AUDIT_LOG, XSA_AUDIT_LOG, and ALL_AUDIT_LOG system views.
BACKUP ADMIN	Authorizes BACKUP and RECOVERY statements for defining and initiating backup and recovery procedures. It also authorizes changing system configuration options with respect to backup and recovery.
BACKUP OPERATOR	Authorizes the BACKUP statement to initiate a backup.

System Privilege	Description
CATALOG READ	Authorizes unfiltered access to the data in the system views that a user has already been granted the SELECT privilege on. Normally, the content of these views is filtered based on the privileges of the user. CATALOG READ does not allow a user to view system views on which they have not been granted the SELECT privilege.
CERTIFICATE ADMIN	Authorizes the changing of certificates and certificate collections that are stored in the database.
CLIENT PARAMETER ADMIN	Authorizes a user to override the value of the CLIENT parameter for a database connection or to overwrite the value of the $$client$$ parameter in an SQL query.
CREATE CLIENTSIDE ENCRYPTION KEYPAIR	Authorizes a user to create client-side encryption key pairs.
CREATE R SCRIPT	Authorizes the creation of a procedure by using the language R.
CREATE REMOTE SOURCE	Authorizes the creation of remote data sources by using the CREATE REMOTE SOURCE statement.
CREATE SCENARIO	Controls the creation of calculation scenarios and cubes (calculation database).
CREATE SCHEMA	Authorizes the creation of database schemas using the CREATE SCHEMA statement.
CREATE STRUCTURED PRIVILEGE	Authorizes the creation of structured (analytic privileges). Only the owner of the privilege can further grant or revoke that privilege to other users or roles.
CREDENTIAL ADMIN	Authorizes the use of the statements CREATE CREDENTIAL, ALTER CREDENTIAL, and DROP CREDENTIAL.
DATA ADMIN	Authorizes reading all data in the system views. It also enables execution of Data Definition Language (DDL) statements in the SAP HANA database. A user with this privilege cannot select or change data in stored tables for which they do not have access privileges, but they can drop tables or modify table definitions.
DATABASE ADMIN	Authorizes all statements related to tenant databases, such as CREATE, DROP, ALTER, RENAME, BACKUP, and RECOVERY.
DATABASE START	Authorizes a user to start any database in the system and to select from the M_DATABASES view.

System Privilege	Description
DATABASE STOP	Authorizes a user to stop any database in the system and to select from the M_DATABASES view.
DROP CLIENTSIDE ENCRYPTION KEYPAIR	Authorizes a user to drop other users' client-side encryption key pairs.
ENCRYPTION ROOT KEY ADMIN	Authorizes all statements related to management of root keys: Allows access to the system views pertaining to encryption (for example, ENCRYPTION_ROOT_KEYS, M_ENCRYPTION_OVERVIEW, M_PERSISTENCE_ENCRYPTION_STATUS, M_PERSISTENCE_ENCRYPTION_KEYS, and so on).
EXPORT	Authorizes EXPORT to a file on the SAP HANA server. The user must also have the SELECT privilege on the source tables to be exported.
EXTENDED STORAGE ADMIN	Authorizes the management of SAP HANA dynamic tiering and the creation of extended storage.
IMPORT	Authorizes the import activity in the database using the IMPORT statements. The user must also have the INSERT privilege on the target tables to be imported.
INIFILE ADMIN	Authorizes making changes to system settings.
LDAP ADMIN	Authorizes the use of the CREATE | ALTER | DROP | VALIDATE LDAP PROVIDER statements.
LICENSE ADMIN	Authorizes the use of the SET SYSTEM LICENSE statement to install a new license.
LOG ADMIN	Authorizes the use of the ALTER SYSTEM LOGGING [ON | OFF] statements to enable or disable the log flush mechanism.
MONITOR ADMIN	Authorizes the use of the ALTER SYSTEM statements for events.
OPTIMIZER ADMIN	Authorizes the use of the ALTER SYSTEM statements concerning SQL PLAN CACHE and ALTER SYSTEM UPDATE STATISTICS statements, which influence the behavior of the query optimizer.

System Privilege	Description
RESOURCE ADMIN	Authorizes statements concerning system resources (for example, the ALTER SYSTEM RECLAIM DATAVOLUME and ALTER SYSTEM RESET MONITORING VIEW statements). It also authorizes many of the statements available in the Management Console.
ROLE ADMIN	Authorizes the creation and deletion of roles by using the CREATE ROLE and DROP ROLE statements. It also authorizes the granting and revoking of roles by using the GRANT and REVOKE statements. 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.
SAVEPOINT ADMIN	Authorizes the execution of a savepoint using the ALTER SYSTEM SAVEPOINT statement.
SCENARIO ADMIN	Authorizes all calculation scenario-related activities (including creation).
SERVICE ADMIN	Authorizes the ALTER SYSTEM [START|CANCEL|RECONFIGURE] statements for administering system services of the database.
SESSION ADMIN	Authorizes the ALTER SYSTEM commands concerning sessions to stop or disconnect a user session or to change session variables.
SSL ADMIN	Authorizes the use of the SET...PURPOSE SSL statement. It also allows access to the PSES system view.
STRUCTUREDPRIVILEGE ADMIN	Authorizes the creation, reactivation, and dropping of structured (analytic) privileges.
TENANT ADMIN	Authorizes the tenant operations performed by the ALTER SYSTEM [RESUME|SUSPEND] TENANT statements.
TABLE ADMIN	Authorizes LOAD, UNLOAD and MERGE of tables and table placement.
TRACE ADMIN	Authorizes the use of the ALTER SYSTEM...TRACES statements for operations on database trace files and authorizes changing trace system settings.
TRUST ADMIN	Authorizes the use of statements to update the trust store.

System Privilege	Description
USER ADMIN	Authorizes the creation and modification of users by using the CREATE | ALTER | DROP USER statements.
VERSION ADMIN	Authorizes the use of the ALTER SYSTEM RECLAIM VERSION SPACE statement of the multi-version concurrency control (MVCC) feature.
WORKLOAD ADMIN	Authorizes execution of the workload class and mapping statements (for example, CREATE | ALTER | DROP WORKLOAD CLASS, and CREATE | ALTER | DROP WORKLOAD MAPPING).
WORKLOAD ANALYZE ADMIN	Used by the Analyze Workload, Capture Workload, and Replay Workload applications when performing workload analysis.
WORKLOAD CAPTURE ADMIN	Authorizes access to the monitoring view M_WORKLOAD_CAPTURES to see the current status of capturing and captured workloads, as well of execution of actions with the WORKLOAD_CAPTURE procedure.
WORKLOAD REPLAY ADMIN	Authorizes access to the monitoring views M_WORKLOAD_REPLAY_PREPROCESSES and M_WORKLOAD_REPLAYS to see current status of preprocessing, preprocessed, replaying, and replayed workloads, as well as the execution of actions with the WORKLOAD_REPLAY procedure.
`<identifier>.<identifier>`	Components of the SAP HANA database can create new system privileges. These privileges use the component-name as the first identifier of the system privilege and the component-privilege-name as the second identifier.

Repository System Privileges

> **i Note**
>
> The following privileges authorize actions on individual packages in the SAP HANA repository, used in the SAP HANA Extended Services (SAP HANA XS) classic development model. With SAP HANA XS advanced, source code and web content are no longer versioned and stored in the repository of the SAP HANA database.

System Privilege	Description
REPO.EXPORT	Authorizes the export of delivery units for example

System Privilege	Description
REPO.IMPORT	Authorizes the import of transport archives
REPO.MAINTAIN_DELIVERY_UNITS	Authorizes the maintenance of delivery units (DU), DU vendor and system vendor must be the same
REPO.WORK_IN_FOREIGN_WORK-SPACE	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_CONTRIBU-TION	

Related Information

GRANT Statement (Access Control)
Developer Authorization in the Repository

8.2.4.3 Object Privileges (Reference)

Object privileges are used to allow access to and modification of database objects, such as tables and views.

The following table describes the supported object privileges in an SAP HANA database.

Object Privilege	Command Types	Applies to	Privilege Description
ALL PRIVILEGES	DDL & DML	• Schemas • Tables • Views	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. This privilege collection is dynamically evaluated for the given grantor and object.

Object Privilege	Command Types	Applies to	Privilege Description
ALTER	DDL	• Schemas • Tables • Views • Functions/procedures	Authorizes the ALTER statement for the object.
CREATE ANY	DDL	• Schemas • Tables • Views • Sequences • Functions/procedures • Remote sources • Graph workspaces	Authorizes all CREATE statements for the object.
CREATE OBJECT STRUCTURED PRIVILEGE	DDL	• Schemas • Views	Authorizes creation of structured privilege commands on the object even if the user does not need to have the CREATE STRUCTURED PRIVILEGE.
CREATE VIRTUAL FUNCTION	DDL	• Remote sources	Authorizes creation of virtual functions (the REFERENCES privilege is also required).
CREATE VIRTUAL PROCEDURE	DDL	• Remote sources	Authorizes creation of virtual procedure to create and run procedures on a remote source.
CREATE VIRTUAL PACKAGE	DDL	• Schemas	Authorizes creation of virtual packages that can be run on remote sources.
CREATE VIRTUAL TABLE	DDL	• Remote sources	Authorizes the creation of proxy tables pointing to remote tables from the source entry.
CREATE TEMPORARY TABLE	DDL	• Schemas	Authorizes the creation of a temporary local table, which can be used as input for procedures, even if the user does not have the CREATE ANY privilege for the schema.

Object Privilege	Command Types	Applies to	Privilege Description
DEBUG	DML	• Schemas • Calculation Views • Functions/procedures	Authorizes debug functionality for the procedure or calculation view or for the procedures and calculation views of a schema.
DEBUG MODIFY	DDL	• Functions/procedures	For internal use only.
DELETE	DML	• Schemas • Tables • Views • Functions/procedures	Authorizes the DELETE and TRUNCATE statements for the object. 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).
DROP	DDL	• Schemas • Tables • Views • Sequences • Functions/procedures • Remote sources • Graph workspaces	Authorizes the DROP statements for the object.
EXECUTE	DML	• Schemas • Functions/procedures	Authorizes the execution of a SQLScript function or a database procedure by using the CALLS or CALL statement respectively. It also allows a user to execute a virtual function.
INDEX	DDL	• Schemas • Tables	Authorizes the creation, modification, or dropping of indexes for the object.

Object Privilege	Command Types	Applies to	Privilege Description
INSERT	DML	• Schemas • Tables • Views	Authorizes the INSERT statement for the object. The INSERT and UPDATE privilege are both required on the object to allow the REPLACE and UPSERT statements to be used. While INSERT applies to views, it only applies to updatable views (views that do not use a join, do not contain a UNION, and do not use aggregation).
REFERENCES	DDL	• Schemas • Tables	Authorizes the usage of all tables in this schema or this table in a foreign key definition, or the usage of a personal security environment (PSE). It also allows a user to reference a virtual function package.
SELECT	DML	• Schemas • Tables • Views • Sequences • Graph workspaces	Authorizes the SELECT statement for the object or the usage of a sequence. When selection from system-versioned tables, users must have SELECT on both the table and its associated history table.
SELECT CDS METADATA	DML	• Schemas • Tables	Authorizes access to CDS metadata from the catalog.
SELECT METADATA	DML	• Schemas • Tables	Authorizes access to the complete metadata of all objects in a schema (including procedure and view definitions), including objects that may be located in other schemas.

Object Privilege	Command Types	Applies to	Privilege Description
TRIGGER	DDL	• Schemas • Tables	Authorizes the CREATE TRIGGER/DROP TRIGGER statement for the specified table or the tables in the specified schema.
UNMASKED	DML	• Schemas • Views • Tables	Authorizes access to masked data in user-defined views and tables. This privilege is required to view the original data in views and tables that are defined by using the WITH MASK clause.
UPDATE	DML	• Schemas • Tables • Views	While UPDATE applies to views, it only applies to updatable views (views that do not use a join, do not contain a UNION, and do not use aggregation).
USERGROUP OPERATOR	DML	• User groups	Authorizes a user to change the settings for a user group, and to add and remove users to/from a user group. Users with the USERGROUP OPERATOR privilege can also create and drop users, but only within the user group they have the USERGROUP OPERATOR privilege on (CREATE USER `<user_name>` SET USERGROUP `<usergroup_name>`). A user can have the USERGROUP OPERATOR privilege on more than one user group, and a user group can have more than one user with the USERGROUP OPERATOR privilege on it.

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.

Related Information

GRANT Statement (Access Control)

8.2.5 Change a Role

You can change the roles and privileges assigned to a role on the *Role Management* page of SAP HANA cockpit.

Prerequisites

- You have the system privilege ROLE ADMIN or are the owner of the role.
- You have the privileges required to grant privileges and roles to the role. For more information, see *Prerequisites for Granting and Revoking Privileges and Roles*.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

It is possible to change a catalog role by revoking roles and privileges from the role or granting further roles and privileges.

> ⚠ Caution
>
> Do not change roles that were originally created in design time, that is, **HDI roles** or **repository roles**. If you change the runtime version of such a role, your changes will be overwritten the next time a new version of the design-time role is deployed. For more information about creating roles in design time, see the SAP HANA developer documentation.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *Role Management*.

   An overview of all existing database roles is displayed.
2. Find the role you want to change.

   > → Tip
   >
   > Search for a role by entering the name or part of the name in the search box.
3. To change the role's LDAP group mapping or comment, choose the *Edit* button in the header area.
4. To change the roles or privileges assigned to the role, select the relevant tab page and choose *Edit*.
5. Make the required changes and save.

Related Information

Database Role Details [page 319]
Prerequisites for Granting and Revoking Privileges and Roles [page 325]

8.2.6 Delete a Role

You can delete a role on the *Role Management* page of the SAP HANA cockpit.

Prerequisites

- You have the system privilege ROLE ADMIN.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *Role Management*.

   An overview of all existing database roles is displayed.
2. Find the role you want to change.

> → Tip
>
> Search for a role by entering the name or part of the name in the search box.

3. Choose *Delete*.

Results

The role is deleted.

> ℹ Note
>
> You can also use the above procedure to delete HDI and repository roles. However, these roles will be recreated when they are deployed again.

8.2.7 Create a User Group

To manage associated users together, you can create and configure a user group.

Prerequisites

- You have the system privilege USER ADMIN.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

User groups are an efficient way to manage users:

- Every user group can have its own dedicated administrator(s).
  In this way, user management tasks can be delegated to several people independently of each other.
- A user group can be configured for exclusive administration.
  This means that only the designated group administrator(s) can manage the users in the group. This could be useful, for example, to protect highly-privileged users or technical users from accidental deletion or manipulation.
- Group-specific user configuration is possible.
  By setting user properties at the group level, you can configure related users quickly and differently to users in other groups. For example, you could include the technical users required by connecting applications in their own user group with a customized password policy.
- To save time, you can use a wizard to quickly apply a basic configuration to user groups.
  For more information, see *Configure the Base Setup for User Groups*.

> **i Note**
>
> User groups are intended to support a separation of user management duties. User groups do not control data access. Data access is controlled by a user's authorizations (roles and privileges).

For more information about user groups, see the *SAP HANA Security Guide*.

Procedure

1. On the *Database Overview*, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *User Group Management*.
2. Choose *New User Group*.
3. Specify the name of the user group.
4. Specify how the user group is to be managed.

Option	Description
Group administration mode:	*Both group administrators and user administrators can manage this user group* (shared administration) Not only designated group administrators can modify the group, but also any user administrator (that is, any user with system privilege USER ADMIN). This is the default administration mode. *Only a group administrator can manage this user group* (exclusive administration) Only the designated group administrator(s) can modify the group, for example, adding new users to the group or removing users from the group. In this way, groups of users can be managed completely independently of each other. **i Note** To add an existing user from the global pool of users to the user group, as well as remove a user from the group (and return it to the global pool of users), the group administrator also needs to be a user administrator, that is have system privilege USER ADMIN.
(Optional) Group creator can manage group	You, as creator and owner of the group, are **automatically** a group administrator because you have the object privilege USERGROUP OPERATOR on the user group. To prevent yourself from being a group administrator, deselect *Group creator can manage group*. You will then not be granted the object privilege USERGROUP OPERATOR

Option	Description
	on the user group. However, you can still grant it to other users to authorize them to be group administrators.
(Optional) Comment	Enter a comment or text to describe the user group.

5. Choose *Add* to save the user group.

Results

The user group is created.

Next Steps

- Designate the user group administrators.
  You do this by granting the object privilege USERGROUP OPERATOR on the user group to the relevant users or roles.
- Add users to the user group.
  You can do this by choosing *Move Users to Group* or *Create User*, depending on whether you want to add an existing user or a new user.
- If required, configure a group-specific password policy.

> → Remember
>
> If you configured the user group for exclusive administration, only a group administrator can add/remove users and configure the password policy. The group administrator may also need to be a user administrator, that is have system privilege USER ADMIN.

Related Information

User Groups
SQL Statements and Authorization for User Group Administration (Reference) [page 346]
Create a Database User [page 349]
Assign Privileges to a User [page 360]
Configure a Password Policy for a User Group [page 283]
Create a Workload Class [page 460]
Create a Workload Class Mapping [page 462]
Configure the Base Setup for User Groups [page 343]
SAP HANA Security Guide
SAP HANA Administration Guide

8.2.7.1 Configure the Base Setup for User Groups

You can use a wizard to quickly apply a basic configuration for user groups.

Prerequisites

- You have the system privilege USER ADMIN.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Instead of manually configuring user groups, you can use a wizard to apply SAP's recommended configuration settings. This allows you to quickly start working with user groups.

For a registered database, the UI notifies you if user group configuration was not completed using the base setup wizards. You have the option to disable these notifications.

> **i** Note
>
> The base setup does not guarantee that the configuration is optimal for your particular system. To optimize the configuration, it may be advisable to manually configure specific settings.

Procedure

1. On the *Database Overview*, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *User Group Management*.
2. Choose *Base Setup*.

   Specify the base settings for the user group.

Option	Description
Display base setup notifications for User Group Management	Enable or disable notifications for a user group type. You can change this setting at any time.
Change Global Settings	To enable or disable all security-related notifications about the base setup, choose *User Interface Notification* > *Display base setup notifications*. You can enable or disable notifications at any time.

Option	Description
	For more information about SAP HANA cockpit user settings, see *Personalizing the Database Overview Page*.
User groups	To complete the base setup, choose *Complete Setup* for each user group type. Follow the on-screen instructions. If you do not wish to change the configuration, choose *Set manually to completed* and choose *OK*. This disables notifications about the base setup.

You can configure the following settings for user groups:

- The user group name
- The group administration mode
- Password policy.
  For example, the permitted number of failed login attempts, when to enforce a password change.
  The field border color alerts you to the status of a specified value:
  - **Blue**: SAP-recommended value
  - **Yellow**: Different from the SAP-recommended value
  - **Red**: Not permitted

3. Choose *Review* to display an overview of the user group settings.

   To make changes, choose *Edit*.
4. Choose *Save*.

Results

The configuration settings are saved.

Notifications about the base setup are no longer displayed.

Related Information

Personalize the SAP HANA Cockpit [page 19]
Create a User Group [page 340]

8.2.7.2 User Group Details

On the *User Group Management* page of the SAP HANA cockpit, you can view the details of a user group.

Field	Description
Owner	User who created the group
Group Administration Mode	Specifies who can manage the group: • Group administrators **and** user administrators (shared administration) Not only designated group administrators can modify the group but also any user administrator (that is, any user with system privilege USER ADMIN). This is the default administration mode. • Group administrators only (exclusive administration) Only the designated group administrator(s) can modify the group, for example, adding new users to the group or removing users from the group. In this way, groups of users can be managed completely independently of each other. **i Note** To add an existing user from the global pool of users to the user group, as well as remove a user from the group (and return it to the global pool of users), the group administrator also needs to be a user administrator, that is have system privilege USER ADMIN. The user administrator can configure the administration mode when creating the group, or later by editing the group.
Comment	Free-text comment or description
Users in Group	Database users in the group A user can belong to only one user group. **i Note** Users don't have to be a member of any group. Users who are not in any group are managed as normal by user administrators.

Field	Description
Password Policy	Group-specific configuration of the password policy The user group password policy is effective if a group-specific value has been configured for at least one option and the group policy has been enabled. For those options with no group-specific value, the value from the database password policy is copied. If no group-specfic values have been configured for any parameters or the group policy has not been enabled, the password policy configured for the database is the effective.
Workload Classes	If a workload class is assigned to the user group, it is displayed. From here, you can: • Display more information about a workload class. Click its name. • Change the configuration of the workload class. For example, to restrict resource usage- Choose *Edit*. • Create a new workload class and assign it to the user group. Choose *Create*, and specify the configuration. Choose *Add User Group*, then *Create*. For more information, see *Managing Workload Classes* in the *SAP HANA Administration Guide*.

Related Information

Managing Workload Classes [page 457]

8.2.7.3 SQL Statements and Authorization for User Group Administration (Reference)

Creating and configuring user groups, and subsequently managing the users in those groups requires different combinations of privileges.

Creating and Configuring User Groups

User administrators create and configure user groups. Group administrators can change the configuration.

To...	You need...	For example...
Create a user group	System privilege USER ADMIN	• `CREATE USERGROUP TechnicalUsers;` • `CREATE USERGROUP Research DISABLE USER ADMIN;`
Change the administration mode of a user group	System privilege USER ADMIN	• `ALTER USERGROUP TechnicalUsers DISABLE USER ADMIN;` • `ALTER USERGROUP Research ENABLE USER ADMIN;`
Make another user the group administrator of a user group	• System privilege USER ADMIN, or • Object privilege USERGROUP OPERATOR on the group with the option to grant it to others	`GRANT USERGROUP OPERATOR ON USERGROUP TechnicalUsers TO TechnicalUsersAdmin;`
Configure user parameters and enable/disable usergroup parameter sets	• System privilege USER ADMIN, or • Object privilege USERGROUP OPERATOR on the group **i Note** If the user group has been configured for exclusive administration, USERGROUP OPERATOR on the group is required.	• `CREATE USERGROUP Training SET PARAMETER 'password_layout' = 'A1a!', 'minimal_password_length' = '7' ENABLE PARAMETER SET 'password policy';` • `ALTER USERGROUP TechnicalUsers SET PARAMETER 'force_first_password_change' = 'false';` • `ALTER USERGROUP TechnicalUsers ENABLE PARAMETER SET 'password policy';` • `ALTER USERGROUP TechnicalUsers DISABLE PARAMETER SET 'password policy';`
Delete a user group	System privilege USER ADMIN **i Note** You cannot delete a user group if there are still users in the group.	`DROP USERGROUP Training;`

Managing Users

Managing users who are not in any user group

User administrators manage users who do not belong to a user group.

To...	You need...	For example...
Create, change, delete a user not in any user group	System privilege USER ADMIN	`CREATE USER Michael PASSWORD <password>;`

Managing users in a user group configured for shared administration (default)

Group administrators and user administrators can manage users in user groups configured with the option `ENABLE USER ADMIN`.

To..	You need...	For example...
Create a user in the user group	• System privilege USER ADMIN, or • Object privilege USERGROUP OPERATOR on the group	`CREATE USER John PASSWORD <password> SET USERGROUP Research;`
Delete a user in the user group	• System privilege USER ADMIN, or • Object privilege USERGROUP OPERATOR on the group	`DROP USER John CASCADE;`
Add an existing user from the global pool of users to the user group	System privilege USER ADMIN	`ALTER USER Julie SET USERGROUP Research;`
Move a user from another user group to the user group	• System privilege USER ADMIN, or • Object privilege USERGROUP OPERATOR on both groups	`ALTER USER Julie SET USERGROUP Training;`
Remove a user from the user group (and return to the global pool of users)	System privilege USER ADMIN	`ALTER USER Julie UNSET USERGROUP;`

Users in a user group configured for exclusive administration

Only group administrators can manage users in user groups configured with the option `DISABLE USER ADMIN`.

To..	You need...	For example...
Create a user in the user group	Object privilege USERGROUP OPERATOR on the user group	`CREATE USER sapsid PASSWORD <password> SET USERGROUP TechnicalUsers;`
Delete a user in the user group	Object privilege USERGROUP OPERATOR on the user group	`DROP USER sapsid CASCADE;`
Add an existing user from the global pool of users to the user group	• System privilege USER ADMIN **and** • Object privilege USERGROUP OPERATOR on the group	`ALTER USER Thomas SET USERGROUP TechnicalUsers;`

To..	You need...	For example...
Move a user from another user group to the user group	• Object privilege USERGROUP OPERATOR on both groups **i Note** If the user's current group is not configured for exclusive administration, object privilege USERGROUP OPERATOR on this group is not required; system privilege USER ADMIN is sufficient.	`ALTER USER Julie SET USERGROUP TechnicalUsers;`
Remove a user from the user group	• System privilege USER ADMIN **and** • Object privilege USERGROUP OPERATOR on the group	`ALTER USER Thomas UNSET USERGROUP;`

Related Information

CREATE USERGROUP Statement (Access Control)
ALTER USERGROUP Statement (Access Control)
CREATE USER Statement (Access Control)
ALTER USER Statement (Access Control)
Password Policy Configuration Options

8.2.8 Create a Database User

You create a standard database user for every person who needs to work directly with the SAP HANA database. When you create a user, you also configure how the user will be authenticated and optionally add the user to a user group.

Prerequisites

- You have the system privilege `USER ADMIN`.
- If you are creating the user in a user group configured for exclusive administration, you must have the privilege USERGROUP OPERATOR on the user group. For more information about user groups, see the *SAP HANA Security Guide*
- If you are integrating SAP HANA database users into a single sign-on (SSO) environment using one or more of the supported mechanisms, the necessary infrastructure must be in place and configured. For more information about SSO, see the *SAP HANA Security Guide*.

- If you are implementing LDAP group authorization or LDAP-based authentication, the necessary infrastructure must be in place and configured. For more information, see the section on configuring LDAP group authorization in the *SAP HANA Administration Guide*.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *User Management*.

   An overview of all existing database users is displayed.
2. Create a new user. Click the + *(Add)* icon, then choose *Create User*.
3. Specify the new user name.

   You must give the user a unique name. User names can contain any CESU-8 characters except for a small subset. For more information, see *Unpermitted Characters in User Names*.
4. Optional: Select the user group to which you want to the assign the user.

   > → Remember
   >
   > If the user group is configured for exclusive administration, you must have the privilege USERGROUP OPERATOR on the user group.
5. Optional: Specify the user's e-mail address.
6. Optional: Specify a validity period for the user, including the appropriate time zone.

   For example, if you are creating a user for a new employee, you can enter their start date in the *Valid From* field.

   If you do not enter any values, the user is immediately and indefinitely valid.
7. Optional: Prevent the user from being able to create objects in his own database schema by selecting *No* for the option *Creation of Objects in Own Schema*.

   > ℹ Note
   >
   > If you select *No* for both this option **and** the next option (*PUBLIC Role*), the user will be created as a restricted user, not a standard user.
8. Optional: Prevent the user from being granted the standard PUBLIC role by selecting *No* for the option *PUBLIC Role*.

   The PUBLIC role contains the privileges for filtered read-only access to the system views. To see data in a particular view, the user also needs the SELECT privilege on the view.

   > ℹ Note
   >
   > If you select *No* for both this option **and** the previous option (*Creation of Objects in Own Schema*), the user will be created as a restricted user, not a standard user.
9. Optional: Prevent the user from being able to connect to the database via ODBC and JDBC clients by selecting the corresponding checkbox.

By default, standard users have access via ODBC and JDBC clients.

If you disable ODBC/JDBC client access, the user can still connect via HTTP. Furthermore, disabling ODBC/JDBC access does not affect the user's authorizations or prevent the user from executing SQL commands via channels other than JDBC/ODBC.

10. Optional: Enter a comment or text to describe the user.
11. Optional: Set the authorization mode to LDAP if the user's authorization is based on LDAP group membership.

    A user with authorization mode *LDAP* is granted roles exclusively based on their LDAP group membership. It is not possible to grant such a user other roles or privileges directly.

    The default user authorization mode is *Local*. This means that the user must be granted roles and privileges directly as normal.

    > **i Note**
    >
    > Setting the authorization mode of the user is only one step in the configuration of LDAP group authorization. For more information, see the section on configuring LDAP group authorization in the *SAP HANA Administration Guide*.

12. Specify how the user can be authenticated.

    > **i Note**
    >
    > You must specify at least one authentication mechanism. For more information about the supported mechanisms, see *Database User Details*.

Authentication Mechanism	Required Configuration
User name and password	Select the type of password. The user can be authenticated by a password stored in the SAP HANA database (option: *Local*) or a password stored in an LDAP directory server (option: *LDAP*). If you choose authentication by local password, enter and confirm the user's initial password. You can override the password policy setting (*Password Change Required on First Logon*) that forces users to change a password set by a user administrator the first time they log on. This is useful for technical users, for example.
Kerberos	Enter the user principal name (UPN) specified in the Microsoft Active Directory or the Kerberos Key Distribution Center as the external ID.
SAP logon and assertion tickets	No additional user configuration required in user definition
SAML	Choose *Add Identity Provider*, select the identity provider, and then enter the user ID known to the SAML identity provider. Alternatively, you can allow the identify provider to map its users to the database user by enabling automatic mapping by provider.

Authentication Mechanism	Required Configuration
	i Note The identity provider must already be created. You can do this on the *SAML Identity Provider* page or using the SQL statement `CREATE SAML PROVIDER`.
JWT (JSON Web Token)	Choose *Add Identity Provider*, select the identity provider, and then enter the user ID known to the JWT identity provider. Alternatively, you can allow the identify provider to map its users to the database user by enabling automatic mapping by provider. **i Note** The identity provider must already be created. You can do this on the *JWT Identity Provider* page or using the SQL statement `CREATE JWT PROVIDER`.
X.509 certificate	Choose *Add X509 Certificate Manually* and enter the user's public key certificate information. **i Note** X.509 certificates are supported only for HTTP access through the SAP HANA XS classic server.

13. Optional: Specify additional user properties required by client applications.
    You can select from the available properties (see *Database User Details*) or manually enter a property.
14. Save the user.

Results

The user is created and appears in the list of users on the left. A new schema is created for the user in the catalog. It has the same name as the user.

Next Steps

Assign roles or privileges to the user (authorization mode *Local* only).

Related Information

User Groups

Create a User Group [page 340]
Single Sign-On Integration
Configure LDAP Group Authorization
Unpermitted Characters in User Names [page 353]
Database User Details [page 312]
Assign Roles to Database Users [page 358]
Assign Privileges to a User [page 360]
Database Users [page 317]
Add a SAML Identity Provider in SAP HANA Cockpit [page 365]
Add a JWT Identity Provider in SAP HANA Cockpit [page 367]

8.2.8.1 Unpermitted Characters in User Names

User names can contain any CESU-8 characters except for a small subset.

The following characters are not allowed as user names:

Unicode Character	Character	Name
U+0021	!	Exclamation mark
U+0022	"	Quotation mark
U+0024	$	Dollar sign
U+0025	%	Percent sign
U+0027	'	Apostrophe
U+0028	(	Left parenthesis
U+0029	)	Right parenthesis
U+002A	*	Asterisk
U+002B	+	Plus sign
U+002C	,	Comma
U+002E	.	Full stop
U+002F	/	Solidus
U+003A	:	Colon
U+003B	;	Semicolon
U+003C	<	Less-than sign
U+003D	=	Equals sign
U+003E	>	Greater-than sign
U+003F	?	Question mark
U+0040	@	Commercial at
U+005B	[	Left square bracket

Unicode Character	Character	Name
U+005C	\	Reverse solidus
U+005D	]	Right square bracket
U+005E	^	Circumflex accent
U+0060	`	Grave accent
U+007B	{	Left curly bracket
U+007C	|	Vertical line
U+007D	}	Right curly bracket
U+007E	~	Tilde

8.2.9 Create a Restricted Database User

You create a restricted user for users who access SAP HANA through client applications – full SQL access via an SQL console is not intended. When you create a restricted user, you also configure how the user will be authenticated and optionally add the user to a user group.

Prerequisites

- You have the system privilege `USER ADMIN`.
- If you are creating the user in a user group configured for exclusive administration, you must have the privilege USERGROUP OPERATOR on the user group. For more information about user groups, see the *SAP HANA Security Guide*
- If you are integrating SAP HANA database users into a single sign-on (SSO) environment using one or more of the supported mechanisms, the necessary infrastructure must be in place and configured. For more information about SSO, see the *SAP HANA Security Guide*.
- If you are implementing LDAP group authorization or LDAP-based authentication, the necessary infrastructure must be in place and configured. For more information, see the section on configuring LDAP group authorization in the *SAP HANA Administration Guide*.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *User Management*.

   An overview of all existing database users is displayed.

2. Create a new restricted user. Click the + *(Add)* icon, then choose *Create Restricted User*.

3. Specify the new user name.

   You must give the user a unique name. User names can contain any CESU-8 characters except for a small subset. For more information, see *Unpermitted Characters in User Names*.

4. Optional: Select the user group to which you want to the assign the user.

   > → Remember
   >
   > If the user group is configured for exclusive administration, you must have the privilege USERGROUP OPERATOR on the user group.

5. Optional: Specify the user's e-mail address.
6. Optional: Specify a validity period for the user, including the appropriate time zone.

   For example, if you are creating a user for a new employee, you can enter their start date in the *Valid From* field.

   If you do not enter any values, the user is immediately and indefinitely valid.

7. Optional: Allow the user to create objects in her own database schema by selecting *Yes* for the option *Creation of Objects in Own Schema*.

   > ℹ Note
   >
   > If you select *Yes* for this option, the user will be created as a standard user, not a restricted user.

8. Optional: Grant the user the standard PUBLIC role by selecting *Yes* for the corresponding option.

   The PUBLIC role contains the privileges for filtered read-only access to the system views. To see data in a particular view, the user also needs the SELECT privilege on the view.

   > ℹ Note
   >
   > If you select *Yes* for this option, the user will be created as a standard user, not a restricted user.

9. Optional: Allow the user to connect to the database via ODBC and JDBC clients by deselecting the corresponding checkbox.

   By default, restricted users are only able to connect to the database using HTTP. You must explicitly allow access via ODBC and JDBC clients by changing this setting.

   For full access to ODBC or JDBC functionality, you must grant restricted users the standard role `RESTRICTED_USER_ODBC_ACCESS` or `RESTRICTED_USER_JDBC_ACCESS`. You can do this on the *Assign Roles* page.

10. Optional: Enter a comment or text to describe the user.
11. Optional: Set the authorization mode to LDAP if the user's authorization is based on LDAP group membership.

    A user with authorization mode *LDAP* is granted roles exclusively based on their LDAP group membership. It is not possible to grant such a user other roles or privileges directly.

    The default user authorization mode is *Local*. This means that the user must be granted roles and privileges directly as normal.

> **i Note**
>
> Setting the authorization mode of the user is only one step in the configuration of LDAP group authorization. For more information, see the section on configuring LDAP group authorization in the *SAP HANA Administration Guide*.

12. Specify how the user can be authenticated.

> **i Note**
>
> You must specify at least one authentication mechanism. For more information about the supported mechanisms, see *Database User Details*.

Authentication Mechanism	Required Configuration
User name and password	Select the type of password. The user can be authenticated by a password stored in the SAP HANA database (option: *Local*) or a password stored in an LDAP directory server (option: *LDAP*). If you choose authentication by local password, enter and confirm the user's initial password. You can override the password policy setting (*Password Change Required on First Logon*) that forces users to change a password set by a user administrator the first time they log on. This is useful for technical users, for example.
Kerberos	Enter the user principal name (UPN) specified in the Microsoft Active Directory or the Kerberos Key Distribution Center as the external ID.
SAP logon and assertion tickets	No additional user configuration required in user definition
SAML	Choose *Add Identity Provider*, select the identity provider, and then enter the user ID known to the SAML identity provider. Alternatively, you can allow the identify provider to map its users to the database user by enabling automatic mapping by provider. **i Note** The identity provider must already be created. You can do this on the *SAML Identity Provider* page or using the SQL statement `CREATE SAML PROVIDER`.
JWT (JSON Web Token)	Choose *Add Identity Provider*, select the identity provider, and then enter the user ID known to the JWT identity provider. Alternatively, you can allow the identify provider to map its users to the database user by enabling automatic mapping by provider.

Authentication Mechanism	Required Configuration
	i Note The identity provider must already be created. You can do this on the *JWT Identity Provider* page or using the SQL statement `CREATE JWT PROVIDER`.
X.509 certificate	Choose *Add X509 Certificate Manually* and enter the user's public key certificate information. **i Note** X.509 certificates are supported only for HTTP access through the SAP HANA XS classic server.

13. Optional: Specify additional user properties required by client applications.
    You can select from the available properties (see *Database User Details*) or manually enter a property.
14. Save the user.

Results

The user is created and appears in the list of users on the left. A new schema is also created for the user in the catalog. It has the same name as the user. However, as a restricted user, the user is not authorized to create objects in this schema. For more information about all restrictions, see *Database Users*.

Next Steps

Assign roles to the user (authorization mode *Local* only).

Related Information

User Groups
Single Sign-On Integration
LDAP Group Authorization
Configure LDAP Group Authorization
Unpermitted Characters in User Names [page 353]
Database User Details [page 312]
Assign Roles to Database Users [page 358]
Database Users [page 317]
Add a SAML Identity Provider in SAP HANA Cockpit [page 365]
SAP HANA Security Guide
SAP HANA Administration Guide

8.2.10 Assign Roles to Database Users

Roles are the preferred way to assign privileges to SAP HANA database users. It is recommended that you assign roles to users instead of granting privileges individually. You can grant (or remove) roles on the *Role Assignment* page of SAP HANA cockpit.

Prerequisites

- The roles you want to assign are available.
- You have the privileges required to grant or revoke roles.
  If you have the system privilege `ROLE ADMIN`, you can grant any role. Otherwise, the following applies:
  - To grant a catalog role, that is a role created in runtime using SQL, you need to have the role being granted yourself and be authorized to grant it to other users and roles, or the system privilege `ROLE ADMIN`.
  - To grant a HDI role, that is a schema-specific role created using the SAP Web IDE and deployed using SAP HANA deployment infrastructure, you need privileges to execute `GRANT_CONTAINER_SCHEMA_ROLES` in the container's API schema, or, if you are a container group administrator, privileges to execute `GRANT_CONTAINER_SCHEMA_ROLES` in the container group's API schema. For more information, see *Prerequisites for Granting and Revoking Privileges and Roles*.
  - To grant a repository role, that is a role created in the repository of the SAP HANA database, you need the object privilege `EXECUTE` on the procedure `GRANT_ACTIVATED_ROLE`.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, choose *Role Assignment* on the *User & Role Management* card..
2. On the *Role Assignment* page:

Option	Description
For a database user: Assign (or remove) one or more roles	Select *User*. Locate a user. An overview of the roles that are assigned to the selected user is displayed. This information includes the user that assigned the roles, and whether the roles can be granted to other users.
For a database role: Assign a role to one or more database users (or remove a role)	Select *Role*. Locate a role.

Option	Description
	An overview of the users that are assigned the role is displayed. This information includes the user group, the user that assigned the role, and whether the roles can be granted to other users.

3. To make changes, choose *Edit*.
4. Depending on whether you are starting from a user or a role, do the following:

Option	Description
For a database user: Assign (or remove) one or more roles	To **assign** one or more new roles, choose *Add.*, select the roles to assign, then choose *OK*. Optional: Authorize the user to grant the role to other users and roles, by selecting *Grantable to Others*. **i Note** This option is not available for roles created in the repository. To **remove** one or more roles, select the roles, then choose *Remove*.
For a database role: Assign a role to one or more database users (or remove a role)	To **assign** the role to one or more users, choose *Assign User*, select the user(s) to which to assign the role, then choose *Assign*. Optional: Authorize the user to grant the role to other users and roles, by selecting *Grantable to Others*. **i Note** This option is not available for roles created in the repository. To **remove** the role from one or more users, select the users, then choose *Remove User*.

i Note

If you are granting an HDI role and you have **both** authorization to execute the `GRANT_CONTAINER_SCHEMA_ROLES` procedure and the system privilege `ROLE ADMIN`, you can choose which granting mechanism to use. Selecting the *HDI Container API* option grants the role by the execution of the procedure. Otherwise, the role is granted through the execution of the GRANT statement.

5. Save.

   The roles are assigned (or removed).

Related Information

Database Roles [page 322]
Prerequisites for Granting and Revoking Privileges and Roles [page 325]

8.2.11 Assign Privileges to a User

It is recommended that you assign roles to users instead of granting privileges individually. However, you can still grant privileges directly to users on the *Privilege Assignment* page.

Prerequisites

- You have the system privilege USER ADMIN.
- You have the privileges required to grant specific privileges to the user.
  To grant SQL privileges, you must have the privilege and/or role yourself and be authorized to grant it to someone else. To grant privileges on activated repository objects, you must be authorized to execute certain stored procedures. For more information, see *Prerequisites for Granting and Revoking Privileges and Roles*.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to the *User & Role Management* area, then choose *Privilege Assignment*.
2. On the *Assign Privileges* page, locate the user you want to edit.

   The user's existing privileges are displayed, as well as the user that assigned them.
3. Assign the required privileges to the user:
   a. For the relevant privilege type, choose *Edit*.
   b. Choose *Add*, then select the privilege(s) you want to assign.

      > **i Note**
      >
      > For object and package privileges, you must first add the object or package and then add the required privilege to the object or package.

   c. Choose *OK*.
   d. Optional: If you want the user to be able to grant the new privilege on to other users, choose *Grantable to Others*.
   e. Save the privilege assignment.

f. Repeat for further privilege types.

Related Information

Prerequisites for Granting and Revoking Privileges and Roles [page 325]
Privileges
System Privileges (Reference) [page 328]
Object Privileges (Reference) [page 333]
Assign Roles to Database Users [page 358]

8.2.12 Change a Database User

For an existing database user, you can change properties such as e-mail, user group, period of validity, or authorization.

Prerequisites

- You have the system privilege `USER ADMIN`.
- If the user is in a user group configured for exclusive administration, you must have the privilege USERGROUP OPERATOR on the user group. For more information about user groups, see the *SAP HANA Security Guide*
- If you are integrating SAP HANA database users into a single sign-on (SSO) environment using one or more of the supported mechanisms, the necessary infrastructure must be in place and configured. For more information about single sign-on integration, see the *SAP HANA Security Guide*.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *User Management*.

   An overview of all existing database users is displayed.
2. Locate the user you want to change.

   > → Tip
   >
   > Search for a user by entering the name or part of the name in the search box, or create a filter by clicking the ▽ *(Filter)* icon.

3. Choose *Edit*.
4. Make the required changes.

   For more information about the individual fields and settings, see *Database User Details*.

   > → Remember
   >
   > To change the user's authorization, open the *Role Assignment* or *Privilege Assignment* pages.

5. *Save*.

Related Information

Database User Details [page 312]
Assign Roles to Database Users [page 358]
Assign Privileges to a User [page 360]
Single Sign-On Integration
SAP HANA Security Guide

8.2.13 Deactivate a Database User

Users can be automatically deactivated for security reasons, for example, if they violate password policy rules. As a user administrator, you may need to explicitly deactivate a user, for example, if an employee temporarily leaves the company or a security violation is detected.

Prerequisites

- You have the system privilege `USER ADMIN`.
- If the user is in a user group configured for exclusive administration, you must have the privilege USERGROUP OPERATOR on the user group. For more information about user groups, see the *SAP HANA Security Guide*

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

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 (Allow only administration users to work on HANA database)*.

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *User Management*.

   An overview of all existing database users is displayed.
2. Locate the user you want to deactivate.

   → Tip

   Search for a user by entering the name or part of the name in the search box, or create a filter by clicking the ▽ *(Filter)* icon.
3. Choose *Deactivate*.

Results

The database user is now deactivated, and remains deactivated until you reactivate it. The user still exists in the database, but cannot connect to the database any more.

ⓘ Note

It may still appear as though deactivated users are still active in the system (for example, when a procedure that was created by the user with DEFINER MODE is called).

To activate the user again, choose *Activate*.

Related Information

SAP Note 1986645
SAP HANA Security Guide

8.2.14 Delete a Database User

You may need to delete a database user, for example, if an employee leaves your organization.

Prerequisites

- You have the system privilege `USER ADMIN`.

- If the user is in a user group configured for exclusive administration, you must have the privilege USERGROUP OPERATOR on the user group. For more information about user groups, see the *SAP HANA Security Guide*

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *User & Role Management*, then choose *User Management*.

   An overview of all existing database users is displayed.
2. Locate the user you want to delete.

   > → Tip
   >
   > Search for a user by entering the name or part of the name in the search box, or create a filter by clicking the ▽ *(Filter)* icon.
3. Choose *Delete*.
4. Specify whether or not you also want to delete dependent objects.

Option	Description
Restrict	Which objects are deleted depends on which objects the user owns: ○ If all objects owned by the user are in the user's home schema: The user, the schema, and all objects in the schema will be deleted. ○ If the user is the owner of any database objects in schemas outside their home schema: The user will **not** be deleted.
Cascade	Delete the user and also all objects owned by the user. ⚠ Caution If you choose *Cascade*: ○ All objects owned by the user will be deleted ○ Privileges granted to others by the user will be revoked ○ All objects in the user's schema will be deleted, **even if they are owned by a different user** All privileges on these objects will also be revoked.

Roles: Local roles (in a user's schema) are always deleted when the schema is deleted.

Other (global) roles do not have an owner. When a user is deleted, users and other roles created by the user are NOT deleted. These global roles can only be deleted explicitly using the DROP ROLE statement.

5. Confirm.

Results

The user is deleted.

Related Information

SAP HANA Security Guide

8.2.15 Add a SAML Identity Provider in SAP HANA Cockpit

If you are implementing Security Assertion Markup Language (`SAML`) to authenticate users accessing SAP HANA via the SQL interface directly (that is, using JDBC and ODBC clients), you must add the SAML identity providers for the required users.

Prerequisites

- You have created a certificate collection with the purpose `SAML` in the database, and have imported the X.509 certificates that will be used to sign the `SAML` assertions from the identity provider. Ensure that the entire certificate chain of the X.509 certificate is available.

  > ⚠ Caution
  >
  > We recommend creating certificate collections for individual purposes in the database directly, rather than using trust stores (PSE) in the file system. By default, the same PSE in the file system is shared by all databases for all external communication channels (including HTTP) and certificate-based authentication. Different PSEs must be explicitly configured for tenant databases.

- You have the system privilege USER ADMIN.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

> **Note**
>
> While you can configure SAML providers for ODBC/JDBC-based SAML authentication using the SAP HANA cockpit, SAP HANA studio or SQL, always use the SAP HANA XS Administration Tool to configure SAML providers that will be used for **HTTP access via the XS classic server**.

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *Security Related Links* and choose *SAML Identity Providers*.
2. Choose *Add Identity Provider*.
   a. Specify a name for the identity provider.

      The following naming conventions apply:

      - Spaces and special characters except underscore (_) are not permitted.
      - The name must start with a letter.
      - The name cannot exceed 127 characters.

   b. Enter the entity ID.
   c. Select the appropriate certificate from the certificate store.

      > **Note**
      >
      > It is not possible to enter the issuer and subject distinguished names (DNs) manually. If the certificate is not available, choose *Go to Certificate Store* and import it. Then, return to the *SAML Identity Providers* page and start again. For more information, see the section on importing a trusted certificate into the certificate store.

   d. Choose *Add*.

Results

The identity provider can now be mapped to individual database users. You can map the identity provider when you create the database user. Alternatively, if the database user already exists, you can change their authentication details.

Related Information

Managing Client Certificates [page 284]
Create a Database User [page 349]
Change a Database User [page 361]
Single Sign-On Using SAML 2.0

8.2.16 Add a JWT Identity Provider in SAP HANA Cockpit

If you are using JSON Web Tokens (JWT) to authenticate users accessing SAP HANA via the SQL interface directly (that is, using JDBC and ODBC clients) or clients that connect to SAP HANA through the SAP HANA XS advanced server, you must add the JWT identity providers for the required users.

Prerequisites

- You have created a certificate collection with the purpose `JWT` in the database and have imported the X.509 certificates that will be used to validate incoming tokens from the identity provider. Ensure that the entire certificate chain of the X.509 certificate is available.
- You have the system privilege USER ADMIN.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Security and User Management* or *All* view selected, navigate to *Security Related Links* and choose *JWT Identity Providers* .
2. Choose *Add Identity Provider*.
   a. Specify a name for the identity provider.
   b. Enter the issuer URL.

      This is used to map the token to an identity provider. It corresponds to the name provided in the `iss` claim of the JWT tokens issued by this JWT provider
   c. Enter the JWT identity claim.

      This is the claim in the JWT token used for mapping the SAP HANA user to an external user name.
   d. Choose *Add*.

Results

The identity provider can now be mapped to individual database users. You map the identity provider when you create the database user. Alternatively, if the database user already exists, you can change the user's authentication details.

Related Information

Managing Client Certificates [page 284]

Create a Database User [page 349]
Change a Database User [page 361]
Single Sign-On Using JSON Web Tokens

8.2.17 Resolve Object Authorization Errors

You can use the Authorization Dependency Viewer in SAP HANA cockpit to analyze database objects and their dependencies.

Prerequisites

- You have the system privilege CATALOG READ.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

The Authorization Dependency Viewer is a graphical tool that shows the object dependency chains of stored procedures and calculation views together with the SQL authorization status of the object owner. The Authorization Dependency Viewer shows you which privileges are missing. This can be a first step to resolving issues with authorizations, invalid objects for stored procedures, or calculation views with complex dependency structures, that are preventing operations from being performed.

You can use the Authorization Dependency Viewer to analyze the following object types:

- NOT AUTHORIZED (258)
- INVALIDATED VIEW (391)
- INVALIDATED PROCEDURE (430)

An object owner must have both the SQL object privilege (for example, EXECUTE, SELECT) and the authorization to grant the object privilege to others (that is, WITH GRANT OPTION is set). If the object owner does not have all the required privileges on all dependent objects, authorization or invalid object errors will occur.

> **i Note**
>
> Use the Authorization Dependency Viewer only with procedures with security mode DEFINER.

> **⚠ Caution**
>
> Grant missing privileges with due care.

Procedure

1. From the *Database Overview*, with the *Security and User Management* or *All* view selected, go to *User & Role Management*.
2. Choose *Authorization Dependency Viewer*.
3. Locate an object to analyze.

   Alternatively, search for a schema and an object name.

   A graphic is displayed showing the relationships between the database objects.

   Each node in the graphic represents a database object. If the same database object is referenced at different levels of the hierarchy, that object may appear multiple times.

   You can search for specific nodes or connecting lines, hide specific objects, and display only objects that are in a line.

   If necessary, use the display options to change the view to focus your analysis.
4. You can display the following information:

To Display...	Do the Following...
The object type and security capabilities	Choose *Show Additional Information*. The following security capabilities can be supported: ○ masking ○ anonymization ○ structured privileges ○ SQL privileges ○ XML privileges You can see whether or not the capabilities are active. If a capability is not supported, no information is displayed for it.
Basic information about the dependency between two nodes	Click the line connecting the two nodes. The lines connecting the nodes indicate the nature and status of the authorization dependency between the objects. If a line is missing, the view is invalid. You can click a connecting line to display the following information: ○ Missing and needed privileges ○ The privilege is grantable or non-grantable

To Display...	Do the Following...
Detailed information about the privileges between two nodes	Go to the *Full Authorization* tab. The graphic displays all the objects accessed. To display the required and missing privileges, click the line connecting the two nodes. If this information is not relevant for the authorization analysis or is internal to SAP HANA, the tab is not available.
Only errors	Go to the *Error Path* tab.

5. In this way, you can isolate the object(s) with missing authorization.
6. Grant the missing privilege(s) to the user with the invalid dependency.

   If you are the object owner, this may be your user, but it may also be the owner of another object.
7. To verify the validity of previously invalid dependencies, refresh the view.

Related Information

SAP Note 1809199

9 Administration View

This view helps you to administer and monitor your database.

The following list describes some of the tasks you can perform on the *Administration* view of your *Database Overview* page:

- View database information and manage database configuration
- Manage workloads
- Manage SQL performance
- View and configure alerts and recommendations
- Work with Smart Data Access

9.1 Managing Tenants from an SAP HANA System Database

Administer and monitor tenant databases from your system database in SAP HANA cockpit.

> **i** Note
>
> Any systems running version SAP HANA 2.0 SP 01, or later runs in multiple-container mode, by default. The cockpit can also monitor single-container systems running earlier versions of SAP HANA. If your system was updated or installed in single-container mode, you can keep it in its current state, but to have tenant databases you must convert it to multiple-container mode.

As the administrator of a tenant (database) system, you are responsible for creating and configuring new tenant databases, subsequently monitoring the availability and performance of databases, and performing certain database administration tasks.

Related Information

Monitoring Tenant Databases in SAP HANA Cockpit [page 418]
Create a Tenant Database [page 381]
Start a Tenant Database [page 383]
Delete a Tenant Database [page 387]

9.1.1 Assign the OS User and Group for High Isolation

Specify the appropriate operating system (OS) user and group when moving a tenant database from a low to a high isolation level.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You are logged in to the system database (SYSTEMDB).
- The operating system user and operating system group exist.

Context

When you change the isolation level to high, the entire system is automatically restarted. Any tenant that doesn't have a specified OS user and OS group isn't able to restart.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select the tenant database to which you want to assign an OS user and group. When the system is running in high isolation mode, any tenant without an OS user and OS group displays a status of Not Running.
3. Click the ... *(More)* icon, and select *Assign OS user & group*.
4. Enter the name of the OS user and OS group you want the tenant to use, and click *OK*.

Next Steps

Start the tenant database.

Related Information

Clear the OS User and Group When Decreasing Isolation [page 378]
Start a Tenant Database [page 383]

9.1.1.1 Increase the System Isolation Level

You can increase the isolation level of an existing system from low (default) to high. With high isolation, the processes of individual tenant databases run under dedicated operating system (OS) users belonging to dedicated (OS) groups and internal communication is secured.

Prerequisites

- You have root access to the SAP HANA system.
- You're logged on to the system database.
- You have the system privilege `DATABASE ADMIN`.
- Internal SAP HANA communication has been appropriately configured for TLS/SSL. `[communication] ssl` in the `global.ini` file must have the value **`false`** (default) or **`systemPKI`**.

  > ⚠ Caution
  >
  > If you're using a manually configured public key infrastructure (PKI) to secure internal communication between hosts, the property `[communication] ssl` must be **`true`**. You can switch to system PKI by changing setting the parameter to **`systemPKI`**.

  For more information, see *Secure Internal Communication* and *Server-Side TLS/SSL Configuration Properties for Internal Communication* in the *SAP HANA Security Guide*.
- If the system is running in an SAP HANA system replication configuration, the system PKI SSFS data file and key file have been copied from the primary system to the same location on the secondary systems:
  - `$DIR_INSTANCE/../global/security/rsecssfs/data/SSFS_<SID>.DAT`
  - `$DIR_INSTANCE/../global/security/rsecssfs/key/SSFS_<SID>.KEY`

Procedure

1. For every tenant database, create a dedicated OS user and group:
   a. As root user, log on to the server on which the name server of the system database is running.
   b. Create new groups for every tenant database:

      ```
      groupadd <groupname>
      ```

   c. Create new users for every tenant database, specifying `sapsys` as the primary group:

      ```
      useradd -g sapsys <username>
      ```

   d. Add every new user to the `sidshm` group and their own group as secondary groups:

      ```
      usermod -G <sid>shm,<usergroup> <username>
      ```

   > ℹ Note
   >
   > If the system is distributed across multiple hosts, you must create identical users and groups on every host. Users and groups must have the same names and IDs on all hosts.

2. Stop all tenant databases in the system.

   In the system database, in an SQL console, execute

   ```
   ALTER SYSTEM STOP DATABASE <databasename>;
   ```

   .

   > → Tip
   >
   > You can also stop tenant databases in the *Database Management* app in 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, in an SQL console, execute:

   ```
   ALTER DATABASE <databasename> OS USER '<username>' OS GROUP '<groupname>'
   ```

   > → Tip
   >
   > You can also assign OS users and groups in the *Database Management* app of the SAP HANA cockpit.

5. Start all tenant databases.

   In the system database, in an SQL console, execute:

   ```
   ALTER SYSTEM START DATABASE <database_name>
   ```

   > → Tip
   >
   > You can also start tenant databases in the *Database Management* app of the SAP HANA cockpit.

Results

The system is now running in high isolation mode. As a result:

- The processes of individual tenant databases run under dedicated OS users belonging to dedicated OS groups, and the processes of the system database run under the `<sid>adm` user.

- Internal database communication is authenticated using X.509 client certificates. Depending on how SSL for internal communication is configured, data communication within databases may also be encrypted. For more information about secure internal communication, see the *SAP HANA Security Guide*.
- Operations that require operating system access are restricted to users with the correct permissions. For more information, see the section on file and directory permissions with high isolation.
- New tenant databases can only be created if a dedicated OS user and group exist.

Related Information

Database Isolation [page 375]
Start a Tenant Database [page 383]
Create a Tenant Database [page 381]
Assign the OS User and Group for High Isolation [page 372]
Secure Internal Communication
File and Directory Permissions with High Isolation [page 377]
Isolation Level High for Backups and Third-Party Backup Tools
Server-Side TLS/SSL Configuration Properties for Internal Communication
SAP HANA Security Guide

9.1.1.1.1 Database Isolation

Every tenant database is self-contained and isolated in terms of users, database catalog, repository, logs, and so on. However, to protect against unauthorized access at the operating system (OS) level, it's possible to increase isolation further through OS user separation and authenticated communication within databases.

OS User Separation

By default, all database processes run under the default OS user `<sid>adm`. If it's important to mitigate against cross-database attacks through OS mechanisms, you can configure the system for high isolation. In this way, the processes of individual tenant databases must run under dedicated OS users belonging to dedicated OS groups, instead of all database processes running under `<sid>adm`. Database-specific data on the file system is then 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's 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.

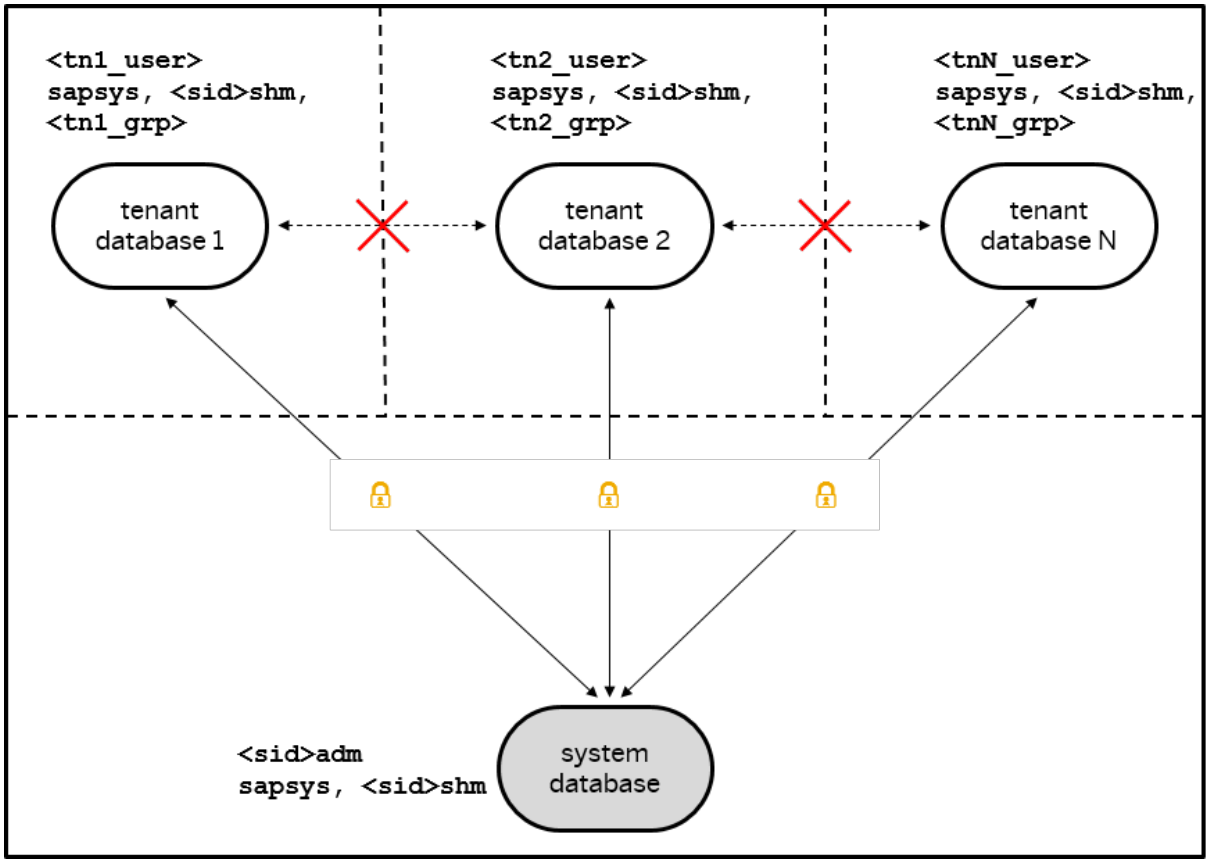

High Database Isolation

Configuration

You can specify the isolation level of the system during installation. The default isolation level is low. It's also possible to change the isolation level of an existing system (from low to high or from high to low) at any time. For more information, see *Increase the System Isolation Level* in the *SAP HANA Administration Guide*. Once high isolation has been configured, a dedicated OS user and group must exist for every tenant database. Otherwise, it's not possible to create or start a tenant database.

Internal database communication is secured with the same mechanism used for securing other internal SAP HANA communication channels. Once high isolation has been configured, authenticated communication within databases is enabled without any change required to the default TLS/SSL configuration for internal communication. However, encryption of data communication may need to be configured explicitly.

Related Information

File and Directory Permissions with High Isolation [page 377]
Secure Internal Communication
Increase the System Isolation Level [page 373]
SAP HANA Administration Guide

9.1.1.1.2 File and Directory Permissions with High Isolation

In an SAP HANA system configured for high isolation, database-specific data on the file system is protected using standard file and directory permissions. All file and directory permissions are managed by the SAP HANA system and don't need to be set by the administrator.

System Database

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

Files and Directories	Tenant OS User in Tenant OS Group	<sid>adm User
Files in directory containing system configuration files	Read permission (644)	Read permission (644)
Files in trace directory of the system database		Read and write permissions (600)
Directory containing Backint parameter file	Read permission (700)	Read permission (700)
Backint parameter file	Read and write permissions (600)	Read and write permissions (600)

Tenant Database

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

> **i** Note
>
> If you want to grant the system administrator access to the tenant database backup files and directories, you need to add the <sid>adm user to each tenant's operating system group.

Files and Directories	Tenant OS User in Tenant OS Group	<sid>adm User
Database-specific directories containing: • Data volumes • Log volumes • Log mirror volumes • Backups	Read, write, and execute permissions (770)	
Database-specific directories containing: • Configuration (*ini) files • Trace files	Read, write, and execute permissions (770)	Read, write, and execute permissions (770)
Files in database-specific directory containing: • Configuration (*ini) files • Trace files	Read and write permissions (666)	Read and write permissions (666)
Directory containing Backint parameter file	Read, write, and execute permissions (750)	Read, write, and execute permissions (750)
Backint parameter file	Read and write permissions (640)	Read and write permissions (640)

Related Information

Working with Third-Party Backup Tools
SAP HANA Security Guide

9.1.2 Clear the OS User and Group When Decreasing Isolation

If you've modified the isolation level from high to low, you can clear the previously assigned operating system user and operating system group.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You are logged in to the system database (SYSTEMDB).

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select the tenant database from which you want to clear the OS user and group.
3. Click the ••• *(More)* icon, and select *Assign OS user & group*.
4. Delete the name of the existing OS user and OS group, and click *OK*.

Next Steps

Start the tenant database.

Related Information

Assign the OS User and Group for High Isolation [page 372]
Start a Tenant Database [page 383]

9.1.2.1 Decrease the System Isolation Level

If you configured a system for high isolation during installation or later, you can decrease it back to the default low level if necessary. With low isolation, the processes of all databases run under the default operating system (OS) user `<sid>adm`.

Prerequisites

- You have root access to the SAP HANA system.
- You're logged on to the system database.
- You have the system privilege DATABASE ADMIN.

Procedure

1. Stop all tenant databases in the system.

   In the system database, execute the SQL statement `ALTER SYSTEM STOP DATABASE <databasename>`.

> → Tip
>
> You can also stop tenant databases in the *Manage Databases* app of the SAP HANA cockpit.

2. Configure the system for low isolation.

   As the operating system user `<sid>adm`, log on to the server on which the master index server is running and run the following command:

   ```
   python /usr/sap/<SID>/HDB<instance>/exe/python_support/convertMDC.py --
   change=databaseIsolation --isolation=low
   ```

   This command runs the following actions:

   - Stops the system
   - Changes the value of the `[multidb] database_isolation` property in the `global.ini` file to `low`
   - Starts the system

3. Clear the assignment of OS users and groups to tenant databases.

   In the system database, in an SQL console, execute

   ```
   ALTER DATABASE <database_name> OS USER '' OS GROUP ''
   ```

   for every tenant database.

   > → Tip
   >
   > You can also clear the OS users and groups in the *Manage Databases* app of the SAP HANA cockpit.

4. Start all tenant databases.

   In the system database, in an SQL console, execute

   ```
   ALTER SYSTEM START DATABASE <database_name>
   ```

   > → Tip
   >
   > You can also start tenant databases in the *Database Management* 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 isn't authenticated.

Related Information

Start a Tenant Database [page 383]
Stop a Tenant Database [page 384]

9.1.3 Create a Tenant Database

You create tenant databases after installation of a multiple-container system, after conversion from a single-container system to a single-tenant system, or anytime a new database is needed

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You are logged in to the system database (SYSTEMDB).
- You have the system privilege `DATABASE ADMIN`.
- If the system is configured for high isolation, the operating system (OS) user and group required for the new tenant database already exist. For more information, see *Database Isolation* in the *SAP HANA Administration Guide*.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select *Create Tenant* > *Create Empty Tenant*.
3. Enter the name of the new database and the password of the SYSTEM user.

   Restrictions apply to tenant names. Alphanumerical string of lowercase alpha characters [a-z], uppercase alpha characters [A-Z], digits [0-9] are permitted. Depending on the file system, tenant names with up to 253 characters are supported.

   > **i** 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. (For high isolation only) Specify the OS user and group of the tenant database.

   If the system in which you're 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. If the system is configured for low isolation (default), all tenant database processes run under the default OS user `<sid>adm`.
5. (Optional) Deselect *Start Automatically* to prevent the database from automatically starting.
6. (For multihost systems only) (Optional) Specify the host on which the database is to be created.

   Under *Advanced Settings*, specify the host you want the master index server to start. If you don't select a host, load-balancing algorithms determine optimal host placement.
7. (Optional) Specify the number of the internal communication port of the master index server.

   Under *Advanced Settings*, specify the port number for the default service. If you don't enter a port, it's assigned automatically based on port number availability. For more information about port number assignment, see *Connections for Tenant Databases* in the *SAP HANA Master Guide*.

8. (Optional) Add additional services required.
   a. In the *Advanced Settings* section, choose *Add Service*.
   b. Select the service you want to add.
   c. (For multihost systems only) (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're automatically determined.
9. Click *Create Tenant Database*. If the host doesn't have at least three available ports, a dialog prompts you with the option to reserve additional instance numbers.

   The system starts creating the database, which can 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 didn't specify host information, load-balancing algorithms determine optimal host placement.
   - If you didn't specify the number of the internal communication port, it's assigned automatically based on port number availability.
   - The necessary data and log volumes are created on the affected hosts.
   - The new database is entered in the M_DATABASES system view of the system database.
   - The `daemon.ini` file is updated and the daemon process is triggered to start the index server 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 *Database Management*. It's now also in the M_DATABASES view (`SELECT * FROM "PUBLIC"."M_DATABASES"`.

Delivery units (DUs) containing automated content start to be deployed in the background. If the system is online, you can monitor the progress of deployment by executing the following statement:

```
SELECT * FROM "PUBLIC"."M_SERVICE_THREADS" WHERE THREAD_TYPE = 'ImportOrUpdate
Content';
```

For more information about automated content, see *SAP HANA Content* in the *SAP HANA Security Guide*.

Next Steps

- Perform a full data backup. For more information, see *Performing Backups* in the *SAP HANA Administration Guide*.
- Adjust the value for the maximum number of asynchronous I/O requests by updating the value of the `fs.aio-max-nr` parameter in `/etc/sysctl.conf`. For more information, see *Linux Kernel Parameters* in the *SAP HANA Administration Guide*.
- Configure the new tenant database as required. For more information, see the section on managing Tenant Databases in the *SAP HANA Administration Guide*.

Related Information

Linux Kernel Parameters
Delete a Tenant Database [page 387]
Start a Tenant Database [page 383]
Monitoring Tenant Databases in SAP HANA Cockpit
CREATE DATABASE Statement (Tenant Database Management)
SAP HANA Administration Guide
SAP HANA Master Guide

9.1.4 Start a Tenant Database

You can start tenant databases either individually, or all at once by starting the whole system.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You are logged in to the system database (SYSTEMDB).
- You have the system privilege DATABASE ADMIN or DATABASE START.

Context

For more information about how to start the whole system, see the sections on stopping and starting a database.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Click the *Start* button for the tenant you want to start.

Related Information

Start a Database [page 164]
ALTER SYSTEM START DATABASE Statement (Tenant Database Management)

M_DATABASES System View
SAP HANA SQL Reference
Prevent the Start of a Tenant Database at System Startup [page 385]

9.1.5 Stop a Tenant Database

You can stop tenant databases either individually, or all at once by stopping the whole system.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You are logged in to the system database (SYSTEMDB).
- You have the system privilege DATABASE ADMIN or DATABASE STOP.

Context

Consider backing up the database before stopping it. For more information about how to stop the whole system, see the sections on stopping and starting a database.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select the tenant database that you want to stop.
3. Click *Stop*.

Results

- The database is stopped.

  > **i** Note
  >
  > This stop is a hard stop, which means that the database is stopped immediately even if users are connected. Open transactions are aborted and rolled back; no savepoint operation is forced. It isn't possible to back up a stopped database.

- The status of database changes accordingly.

Related Information

Stop a Database [page 166]
ALTER SYSTEM STOP DATABASE Statement (Tenant Database Management)
M_DATABASES System View
Creating Backups
SAP HANA SQL Reference

9.1.6 Prevent the Start of a Tenant Database at System Startup

Prevent the start of individual tenant databases.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You are logged in to the system database (SYSTEMDB).
- You have the system privilege DATABASE ADMIN.

Context

By default, all tenant databases that were running before the SAP HANA system was stopped are restarted upon system startup. For troubleshooting purposes, you may want to prevent a particular database from starting until the issue is resolved. You can also prevent startup using the ALTER DATABASE statement.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select the tenant.
3. Click the *Tenant Actions* dropdown menu and select *Set Restart Mode* .
4. Select *No auto-restart*.
5. (Optional) To display the restart mode of all tenants through the ⚙ *(Setting)* icon and select *Restart Mode*.

Results

Alternatively, you can prevent a tenant restart by executing the ALTER DATABASE statement:

```
ALTER DATABASE database_name NO RESTART
```

The tenant database isn't started after a system restart. You can verify the mode by querying the public view M_DATABASES. The result looks like this:

```
| DATABASE_NAME   |DESCRIPTION                        | RESTART_MODE |
|-----------------|-----------------------------------|--------------|
| SYSTEMDB        | SystemDB-<SID>-<INSTANCE>         | DEFAULT      |
| <SID>           | SingleDB-<SID>-<INSTANCE>         | NO           |
```

To restore the default behavior, execute the following ALTER DATABASE statement:

```
ALTER DATABASE <database_name> DEFAULT RESTART
```

Then, start the tenant database manually. At the next system startup, the tenant database will be restarted.

Related Information

Start a Tenant Database [page 383]
Stop a Tenant Database [page 384]
ALTER DATABASE Statement (Tenant Database Management)

9.1.7 Rename a Tenant Database

You can use the cockpit to rename a tenant database.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You are logged in to the system database (SYSTEMDB).
- You have the system privilege DATABASE ADMIN:
- The database to be renamed:
  - Isn't running.
  - Isn't part of a copy or move operation.
  - Doesn't have system replication active.
  - Doesn't have the SAP HANA dynamic tiering option.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select the tenant database that you want to rename.
3. Click *Stop* if the tenant is running.
4. Click the *Tenant Actions* dropdown menu and select *Rename Tenant*.
5. Enter the new name, and select *Rename*.

   Restrictions apply to tenant names. Alphanumerical string of lowercase alpha characters [a-z], uppercase alpha characters [A-Z], digits [0-9], the hyphen (or minus) character "-", and the underscore character "_" are permitted. Depending on the file system, tenant names with up to 253 characters are supported.

Results

The database is renamed. On-disk directories that contain the tenant name are also renamed. Existing backups aren't renamed but backup history remains continuous.

Related Information

Create a Tenant Database [page 381]
Delete a Tenant Database [page 387]
RENAME DATABASE Statement (Tenant Database Management)

9.1.8 Delete a Tenant Database

You can delete tenant databases that are no longer required.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You are logged in to the system database (SYSTEMDB).
- You have the system privilege `DATABASE ADMIN`.

Context

If you delete a tenant database that is running SAP HANA 2.0 SPS 01 or later, you have the option to keep the backup directories of the deleted tenant. Backups can then only be removed by deleting them from the file system. If you delete a tenant database that is running an earlier version of SAP HANA, the backup directories are deleted automatically. It's therefore recommended that if you want to preserve these backup directories, you relocate them before deleting the database.

> **i** Note
>
> If a tenant SAP HANA database is enabled for usage with XS advanced and mapped to an organization/space in XS advanced, then it is recommended not to use cockpit (or the SQL command line interface) to delete the tenant database but to use the XS advanced xs command line interface. This is necessary so that XS advanced is aware that the tenant database is no longer available. Note, too, that if you use the XS advanced xs CLI to delete a tenant database used by XS advanced, all data in the tenant database is lost. See also 'Maintaining Tenant Databases in XS Advanced'.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select the tenant to delete and click *Stop*.

   Once stopped, its status changes to *Stopped*.
3. Click the ⊗ *(Delete)*icon for the tenant.
4. If this database is running SAP HANA 2.0 SPS 01 or later, choose whether to *Keep Backup Directories* or *Delete Directories* and proceed with the database deletion, or *Cancel* the database deletion. If the database is running an earlier version of SAP HANA, choose whether to *Delete Tenant* or *Cancel* the database deletion.

Results

The system commences the process to delete the database. Once deleted, the database disappears from the list. Volumes and trace files are removed.

Next Steps

If you configured the SAP Web Dispatcher to route HTTP requests to the deleted database, you need to update the configuration.

Related Information

ALTER SYSTEM STOP DATABASE Statement (Tenant Database Management)
DROP DATABASE Statement (Tenant Database Management)
Execute SQL Statements in SAP HANA Studio
Configure HTTP(S) Access to Tenant Databases via SAP HANA XS Classic
SAP HANA SQL and System Views Reference
Start a Tenant Database [page 383]

9.1.9 Restrict Features Available to a Tenant Database

To safeguard and/or customize your system, certain features of the SAP HANA database can be disabled in tenant databases.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You are logged in to the system database (SYSTEMDB).
- You have the system privilege INIFILE ADMIN.

Context

Some features of the SAP HANA database aren't required or desirable in certain environments, in particular features that provide direct access to the file system, the network, or other resources. To maximize your control over the security of your system, you can disable these features in tenant databases, for example import and export operations or the ability to back up the database.

The system view M_CUSTOMIZABLE_FUNCTIONALITIES provides information about those features that can be disabled and their status. This view exists in both the SYS schema of every database, where it contains database-specific information, and in the SYS_DATABASES schema of the system database, where it contains information about the enablement of features in all databases.

For more information about the features that can be disabled and why, see *Restricted Features in Tenant Databases* in the *SAP HANA Tenant Databases*.

You can disable features in tenant databases in the `customizable_functionalities` section of the `global.ini` file, as well as in the SAP HANA cockpit as described here.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Click *Restricted Features*.
3. Click the tenant to restrict its features.
4. Select or clear the features to restrict.

   > **i Note**
   >
   > The information in the *Layer* column indicates where the feature was defined. If a feature isn't defined at the database, then system-defined values are used. If there are no system-defined values, then default values are used. Thus, if you choose to restrict a system-defined value, it's restricted on all of the system's tenants.

5. (Optional) To delete a feature, you must be viewing the tenant or system in which it was defined. Select the delete icon at the end of the row.

   > **i Note**
   >
   > If you delete a database-defined feature, the system-defined values are used. If there are no system-defined values, then default values are used.

6. Click *Save* to restrict the features you selected.

Next Steps

Stop and restart the affected tenant database.

Related Information

Stop a Tenant Database [page 384]
Start a Tenant Database [page 383]
Lock Parameters Against Editing for a Tenant Database [page 391]

9.1.9.1 Copy Restricted Features

Copy all features that have been restricted on another tenant and apply them to this tenant.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You are logged in to the system database (SYSTEMDB).
- You have the system privilege INIFILE ADMIN.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Click *Restricted Features*.
3. Select *Copy Restricted Features*.
4. Select the tenant or system from which you wish to copy the restricted features (the source).
5. Select *OK*.

Results

All features that were restricted on the source database or system are now also restricted on this (target) database or system.

9.1.10 Lock Parameters Against Editing for a Tenant Database

To ensure the stability and performance of the overall database or for security reasons, you can prevent certain system parameters from being changed by tenant database administrators, for example, parameters related to database management. A configuration change blacklist is available for this purpose.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

- You are logged in to the system database (SYSTEMDB).
- You have the system privilege INIFILE ADMIN.

Context

Database configuration (*.ini) files have a database layer to facilitate the configuration of system parameters for individual tenant databases. However, it may be desirable to prevent changes to certain parameters being made directly in tenant databases because they could, for example, affect the performance of the database as a whole (CPU and memory management parameters).

You can use the cockpit to blacklist parameters for a particular database—that is, to lock them against editing. (The blacklist is stored in `multidb.ini`.) Several parameters are blacklisted by default, so you see them when you visit the *Blacklisted Parameters for Tenants* page for a tenant. You can remove default properties from that page—that is, make them editable by the tenant—and you can add parameters—lock them so they can't be edited.

> **i** Note
>
> Properties in the blacklist can still be configured at all levels in the system database. For more information about configuring system properties, see *Configuring SAP HANA System Properties (INI Files)*.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. On the *Manage Databases* screen, click *Manage Blacklisted Parameters* (upper right).
3. On the *Blacklisted Parameters for Tenants* page, the active tenant is highlighted in the left pane. Click the name of another tenant to manage its parameters.
4. To add a parameter to the blacklist, click *Add Parameter* at the top of the screen.
5. In the *Add Parameter to Blacklist* dialog, select the configuration file in which the parameter you want to add appears. To add a parameter that appears in more than one file, select *Any with the specified section*.
6. Enter or select your parameter's section in the configuration file. If you start to type the section name, the cockpit offers section names that match what you've entered—click to select one. If you prefer to select the section, click the double box at the end of the *Section* field to see a list. You can even combine the two methods: enter a few characters to narrow the number of choices offered when you click the double box.
7. Enter or select the parameter or parameters you want to add to the blacklist. If you start to type the parameter name, the cockpit offers section names that match what you've entered—click to select one. If you prefer to select the section, click the double box at the end of the *Parameters* field to see a list. You can even combine the two methods: enter a few characters to narrow the number of choices offered when you click the double box.

   You can enter multiple parameters if they're all in the section you specified. You can also specify new parameters.
8. Click *OK* to save the new blacklist entries.

Related Information

Configuring SAP HANA System Properties (INI Files)
Restrict Features Available to a Tenant Database [page 389]

9.1.10.1 Default Blacklisted System Properties in Tenant Databases

In systems that support tenant databases, there's configuration change blacklist `multidb.ini`, which is delivered with a default configuration.

The following table lists the system properties that are included in the `multidb.ini` file by default. So, tenant database administrators can't change these properties. System administrators can still change these properties in the system database in all layers.

You can customize the default configuration change blacklist by changing existing entries in the `multidb.ini` file and adding new ones. For more information about how to prevent changes to specific system properties in tenant databases, see *Prevent Changes to System Properties in Tenant Databases* in the *SAP HANA Administration Guide*.

File/Section	Properties	Description
`auditing configuration`	• `default_audit_trail_type` • `emergency_audit_trail_type` • `alert_audit_trail_type` • `critical_audit_trail_type` • `audit_statement_length`	Prevents configuration of audit trail targets and the maximum audit statement length
`communication`	*	Prevents configuration of default key and trust stores, as well as other critical communication settings
`global.ini/customizable_functionalities`	*	Prevents disabling of restricted features
`global.ini/extended_storage`	*	Prevents configuration of extended storage (SAP HANA dynamic tiering)
`global.ini/persistence`	• `basepath_datavolumes_es` • `basepath_logvolumes_es` • `basepath_databackup_es` • `basepath_logbackup_es`	

File/Section	Properties	Description
global.ini/ system_replication	• keep_old_style_alert • enable_full_sync • operation_mode	Prevents configuration of certain system replication settings
global.ini/ system_replication_communication	*	
global.ini/ system_replication_hostname_resolution	*	
global.ini/xb_messaging	*	Prevents configuration of messaging
multidb.ini/ readonly_parameters	*	Prevents configuration of the multidb.ini file itself
indexserver.ini/ authentication	SapLogonTicketTrustStore	Prevents configuration of the trust store for user authentication with logon/ assertion tickets
memorymanager	• allocationlimit • minallocationlimit • global_allocation_limit • async_free_threshold • async_free_target	Prevents configuration of memory allocation parameters
execution	max_concurrency	Prevents configuration of threading and parallelization parameters
session	• maximum_connections • maximum_external_connections	
sql	sql_executors	

Related Information

Prevent Changes to System Properties in Tenant Databases

Unlock Blacklisted Parameters [page 395]

Copy Blacklisted Parameters [page 396]

9.1.10.2 Unlock Blacklisted Parameters

Remove a parameter from the blacklist for a tenant database so that the parameter can be edited.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the system privilege INIFILE ADMIN.
- You are connected to the system database.

Context

By removing a parameter from the blacklist, you can enable it to be edited. However, you can remove only parameters that you or other users have added to the blacklist—you can't remove parameters that are on the blacklist by default. Default parameters are displayed without delete or edit controls on the *Blacklisted Parameters for Tenants* page.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. On the *Manage Databases* screen, click (upper right).

   To see the *Manage Blacklisted Features* button, you might need to widen the window or click the three dots in the upper right corner of the screen.Manage Blacklisted Parameters
3. On the *Blacklisted Parameters for Tenants* page, the active tenant is highlighted in the left pane. Click the name of another tenant to manage its parameters.
4. To remove a parameter to the blacklist, click the red X to the right of the parameter to be deleted and confirm the deletion.

   The cockpit displays the blacklist without the parameter you removed.

Related Information

Lock Parameters Against Editing for a Tenant Database [page 391]

9.1.10.3 Copy Blacklisted Parameters

Copy parameters that have been added to another tenant's blacklist.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the system privilege INIFILE ADMIN.
- You are connected to the system database.

Context

If you've added parameters to the target tenant before copying, you can choose whether to keep them. Parameters on the blacklist by default aren't copied.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. On the *Manage Databases* screen, click *Manage Blacklisted Parameters* (upper right).
3. On the *Blacklisted Parameters for Tenants* page, the active tenant is highlighted in the left pane. Click the name of another tenant to manage its parameters.
4. Click *Copy Parameters* at the top of the screen.
5. In the *Copy Parameters* dialog, select the tenant whose blacklist you want to copy (the source), and specify the tenant to which you want to copy the parameters (the target).
6. To make this tenant's blacklist an exact copy of the source tenant's blacklist, discarding any parameters you've added, select *Replace all parameters, making target and source identical*. To keep parameters you've added, select *Augment the target parameters with those from the source*.
7. Click *OK* to copy the blacklist.

9.1.11 Create a Fallback Snapshot

Create a fallback snapshot of a tenant database. You can revert the state of a tenant database to a specific point in time if needed.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the privilege `DATABASE ADMIN`.
- You are connected to the system database.

Context

You can create a fallback snapshot for a tenant database. It allows you to revert to a particular database state. If you no longer need the fallback snapshot, you can delete it.

A fallback snapshot can be useful if you perform changes to the contents of a database that you need to roll back quickly, for example, if you upgrade to a new version of an application.

Fallback snapshots can only be created for tenant databases. Configuration changes aren't included. You can only have one fallback snapshot per tenant database. You must delete the existing snapshot before creating a new one. A service can't be added or removed if a fallback snapshot already exists.

A fallback snapshot doesn't replace a database backup. A fallback snapshot isn't included in a database backup. You can create A fallback snapshot when the tenant database is running in the primary system of a system replication setup with the following restrictions:

- Delete any existing fallback snapshots on the primary system before configuring system replication.
- `logreplay` or `logreplay_readaccess`.
- Fallback snapshots are propagated to the secondary system with the continuous log shipping. If a full data shipping is needed to resync the primary and the secondary, the fallback snapshot isn't available on the secondary afterwards.The operation mode for system replication must be

You can also create a fallback snapshot by executing the `ALTER DATABASE` statement in the system database:

```
ALTER DATABASE <database name> CREATE FALLBACK SNAPSHOT;
```

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select the tenant database for which you want to create a fallback snapshot.
3. Select *Fallback Snapshot* > *Create Fallback Snapshot*.

Results

A fallback snapshot is created. You can verify creation by querying the system view SYS_DATABASES.M_SNAPSHOTS.

Related Information

ALTER DATABASE Statement (Tenant Database Management)
M_SNAPSHOTS System View
Reset to a Fallback Snapshot [page 398]
Delete a Fallback Snapshot [page 399]
Operation Modes for SAP HANA System Replication

9.1.12 Reset to a Fallback Snapshot

Reset a tenant database to a fallback snapshot.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the privilege `DATABASE ADMIN`.
- You are connected to the system database.

Context

Resetting can be useful if you performed changes to the contents of a database that you need to roll back quickly.

You can also start a tenant database from a fallback snapshot by executing the `ALTER DATABASE` statement in the system database:

```
ALTER SYSTEM START DATABASE <database name> FROM FALLBACK SNAPSHOT;
```

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select the tenant database that you want to revert to the state captured in the fallback snapshot.
3. Select *Fallback Snapshopt* > *Reset to Fallback Snapshot*.

   The tenant database is reset to the state captured in the fallback snapshot.

Results

The fallback snapshot will remain available after the reset. If you no longer need the fallback snapshot, delete it.

Related Information

Create a Fallback Snapshot [page 397]
Delete a Fallback Snapshot [page 399]

9.1.13 Delete a Fallback Snapshot

Delete a fallback snapshot of a tenant database.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the privilege `DATABASE ADMIN`.
- You are connected to the system database.

Context

You can only create one fallback snapshot per tenant database. If you need to create a new fallback snapshot, delete the existing one first.

You can also delete a fallback snapshot by executing the `ALTER DATABASE` statement in the system database:

```
ALTER DATABASE <database name> DROP FALLBACK SNAPSHOT;
```

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select the tenant database for which you want to delete a fallback snapshot.
3. Select *Fallback Snapshot* > *Delete Fallback Snapshot*.

Related Information

Create a Fallback Snapshot [page 397]
Reset to a Fallback Snapshot [page 398]

9.1.14 Copy or Move a Tenant Database Using Replication

Use system replication to copy or move a tenant database from one system to another.

Prerequisites

- You have configured tenant replication.
- You've navigated to the *Database Overview* page of the system database of the source or target system. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the system privilege `DATABASE ADMIN`.
- You have created a complete system backup for the source tenant database.

  > ⚠ Caution
  >
  > When you use the cockpit to move a tenant, the source database is deleted as part of the process. If the source is running SAP HANA 2.0 SP01 or earlier, its backups are also deleted as part of the process—you can't roll back! Before moving, SAP recommends that you run a backup, then replicate the backup to a new location.

- For the target (the database where you're putting the moved or copied database):
  - You must have the system privileges `DATABASE ADMIN`.
- If the system is configured for high isolation, the operating system (OS) user and group required for the new tenant database already exist. For more information, see *Database Isolation* in the *SAP HANA Administration Guide*.
- You are logged in to the system database (SYSTEMDB).

Context

For more information, see *Copying and Moving Tenant Databases Between Systems* in the *SAP HANA Administration Guide*.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. To copy or move a tenant to another system, choose *Tenant Actions* > *Replicate Tenant*. If you want to create a new tenant based on an existing tenant, choose *Create Tenant* > *Create Tenant Using Replication*. For more information, see *Create a Tenant Database Using Replication*.
3. To create a copy of an existing tenant database, select *Copy using replication*. To remove the original tenant after the copy has been created, select *Move using replication*.
4. Select the source database and tenant.
5. Enter a name for the new tenant.

   Restrictions apply to tenant names. Alphanumerical string of uppercase alpha characters [A-Z], digits [0-9] are permitted, starting with a letter. Depending on the file system, tenant names with up to 253 characters are supported.
6. (Optional) Specify the number of the internal communication port of the listed services.

   Under *Advanced Settings*, specify the port number for each service. If you don't enter a port, it's assigned automatically based on port number availability. In multihost systems enter host and port of a service. For more information about port number assignment, see *Connections for Tenant Databases* in the *SAP HANA Master Guide*.
7. (For high isolation system only) Enter a dedicated OS user and group for the source.
8. If configuration changes are required for the replication, warnings appear to indicate the required changes. Select *Approve* to proceed.

   You may see a warning that you do not have the privilege required to check if the databases are fully configured for tenant replication. Please contact your system administrator to make sure that tenant replication has been properly configured.
9. If prompted, enter the SAP Control credentials required for the restart of the system.
10. If a trust relationship has not yet been established between the source system and the target system, select an existing public key certificate or upload a certificate.
11. Review the summary, then choose *Copy Tenant Database* or *Move Tenant Database* to start replicating the tenant database.

Next Steps

Register the new tenant in SAP HANA Cockpit Manager so it can be managed by SAP HANA Cockpit.

Related Information

Monitoring Tenant Databases in SAP HANA Cockpit [page 418]
Create Data Backups [page 603]
Copying and Moving Tenant Databases

9.1.15 Reset the SYSTEM Password of a Tenant Using the Cockpit

If the password of the SYSTEM user in a tenant database is unknown, you can reset it from the system database.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- There's no user available with the system privilege USER ADMIN that can reset the SYSTEM user password.

  > **i** Note
  >
  > If you can log on as SYSTEM or another user with the system privilege USER ADMIN, don't use the procedure described here to change the password of the SYSTEM user. Instead, change the password using the *User* editor in SAP HANA cockpit

- You're connected to the system database and have the system privilege DATABASE ADMIN.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select the tenant and click *Stop*.

   Once stopped, its status changes to *Not running*.
3. Click *Tenant Actions* and then *Reset SYSTEM Password*.
4. Enter and confirm a new temporary password for the SYSTEM user.
5. Select *Reset Password & Restart*.

Results

- The password for the SYSTEM user is reset and the tenant database is restarted.
- The next time you log on with the SYSTEM user, you will be prompted to change the password in line with the password policy of the tenant database
- If the SYSTEM user was previously deactivated, locked, or expired, it's now activated again. In this case, we recommend that you return it to its deactivated state.
- If auditing is enabled, the password change is automatically logged in both the system and tenant database audit trails.

Related Information

Monitoring Tenant Databases in SAP HANA Cockpit [page 418]
Resetting the SYSTEM User Password

9.1.16 Memory and CPU Usage for Tenant Databases

Manage and control the memory and CPU usage of your system by configuring limits for individual tenant databases. If necessary, you can also reserve memory for the system database.

Managing Resource Usage of Tenant Databases

Several system properties allow you to influence the allocation of memory and CPU resources in SAP HANA systems. System properties (INI) files have a database layer to facilitate the configuration of properties for individual tenant databases.

The following properties are useful for influencing the resource consumption of tenant databases.

- `[memorymanager] allocationlimit` in the `global.ini` file
  Use this property to limit the maximum amount of memory (in MB) that can be allocated individually to processes of a tenant database. Each process of a tenant database can allocate the specified value. Setting the allocation limit too low can 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 happens on the fly and can 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. Don't 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 isn't enough to achieve the desired performance gains, so 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. The system has an average memory consumption of 15 GB.

If the system database is experiencing performance problems, for example, out-of-memory situations, you can reserve a minimum amount of memory (MB) for the system database by configuring the parameter `[multidb] systemdb_reserved_memory` in the `global.ini` file.

Related Information

Controlling Parallel Execution of SQL Statements
Configuring SAP HANA System Properties (INI Files)

9.1.16.1 Define Memory Allocation Limits

As part of the provisioning process, you can ensure that memory is shared appropriately between tenant databases. By setting the memory allocation limits for each database on a host, you ensure appropriate

memory sharing through setting the maximum amount of memory that can be allocated for a particular tenant database.

Prerequisites

-
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You are connected to the SYSTEMDB.

Context

You can use the *Memory Allocation* tab of the *Configure Workload Allocation* app to view and modify memory allocation limits. These limits correspond to the settings of the `allocationlimit` parameter in the memory manager section of the global.ini file.

SAP HANA preallocates and manages its own memory pool, used for storing in-memory table data, thread stacks, temporary results, and other system data structures. When more memory is required for table growth or temporary computations, the SAP HANA memory manager obtains it from the pool. When the pool can't satisfy the request, the memory manager increases the pool size by requesting more memory from the operating system, up to a predefined allocation limit.

You can adjust the allocation limit for the services in the system at the following levels:

- System
- Tenant Database (these settings override the system settings)
- Host (these settings override the database and system settings)

For each database (or each database per host in a multi-host system), you can also refer to a mini-chart representing the memory usage of the indexserver in one day, where the vertical line shows the allocation limit setting for this specific database, dark green shows the actual used memory and light green shows the peak used memory.

The allocation limit value is applied to each service (for example, indexserver, nameserver, compileserver). However, when the *Configure Workload Allocation* app calculates the total memory available for allocation, only the indexserver is considered, since it requires much more memory than the other services.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Select *Configure Workload Allocation* from the overflow menu in the header.
3. Select *Memory Allocation*.
4. View the *Default allocation limit per service*, or choose to edit the default by clicking the pencil icon. Changing the default causes each of the databases and hosts to inherit the value of the allocation limit (except those databases or hosts that have an allocation limit that already differs from the default).

5. View the *Allocation Limit* for each database or host, or choose to edit the value by clicking the pencil.
6. If you choose to edit the limit value, the *Edit Allocation Limit* dialog displays. Enter a value, or clear the allocation limit input field to revert to the allocation limit inherited from the default value.

   The total memory available for allocation is equal to the sum of the allocation limits for each database (regardless of whether you've edited the limit or allowed the default value to be inherited). A warning displays if this sum exceeds the global allocation limit (the amount of available memory per host). However, SAP HANA does allow the sum of the allocation limits to exceed the global allocation limit.

Related Information

SAP HANA Used Memory [page 406]
Memory Sizing [page 408]
Allocated Memory Pools and Allocation Limits [page 409]
SAP HANA Memory Usage and the Operating System [page 410]

9.1.16.1.1 SAP HANA Used Memory

The total amount of memory used by SAP HANA is referred to as used memory. It includes program code and stack, all data and system tables, and the memory required for temporary computations.

SAP HANA consists of a number of processes running in the Linux operating environment. Under Linux, the operating system (OS) is responsible for reserving memory to all processes. When SAP HANA starts up, the OS reserves memory for the program code (sometimes called the text), the program stack, and static data. It then dynamically reserves additional data memory when requested by the SAP HANA memory manager. Dynamically allocated memory consists of heap memory and shared memory.

The following figure shows used memory, consisting of code, stack, and table data:

SAP HANA Used Memory

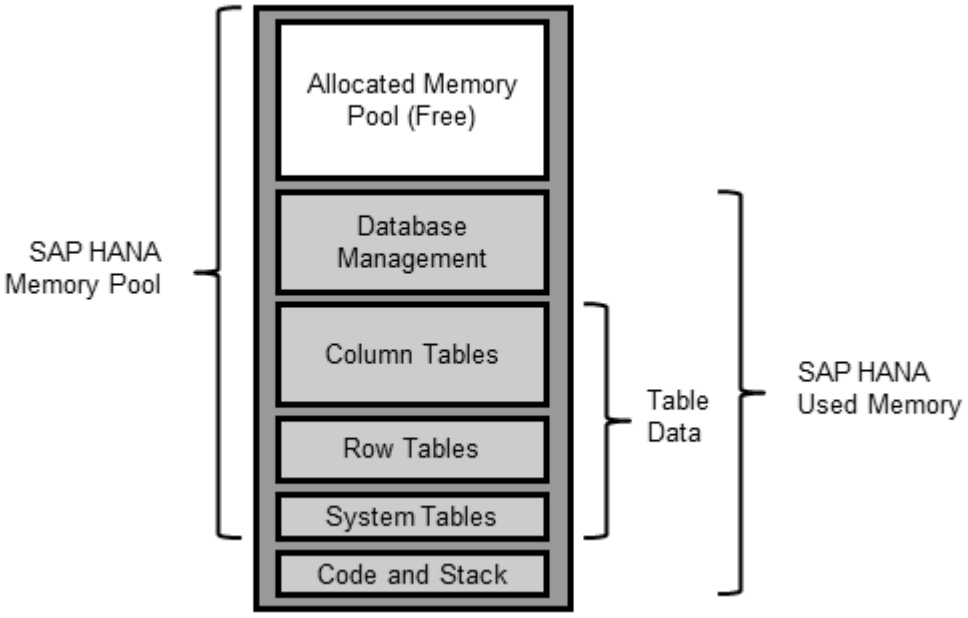

Since the code and program stack size are about 6 GB, almost all of used memory is used for table storage, computations, and database management.

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's 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 track the maximum value for used memory over time.

You can also reset peak used memory. Resetting 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, which is sometimes called "lazy loading". This method means that columns that are unused aren't loaded and memory waste is avoided. When the SAP HANA database runs out of allocatable memory, it tries to free up some memory by unloading unimportant data (such as caches) and even table columns that haven't been used recently. Therefore, if it's important to measure precisely the total, or worst-case, amount of memory used for a particular table, it's important to ensure that you fully load 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, taxes even 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 also shows the peak memory size used to execute expensive statements.

It's 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

The Performance Monitor [page 515]
Monitor Table Usage [page 530]
Load/Unload a Column Table into/from Memory
Monitor and Analyze Active Statements [page 526]
Monitor and Analyze Expensive Statements [page 529]
Setting a Memory Limit for SQL Statements

9.1.16.1.2 Memory Sizing

Memory sizing is the process of estimating in advance the amount of memory that is 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 must 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 applies 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 is required for temporary computations?
  The amount of extra memory depends on the size of the tables (larger tables 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

9.1.16.1.3 Allocated Memory Pools and Allocation Limits

SAP HANA, across its different processes, reserves a pool of memory before actual use. This pool of allocated memory is preallocated from the operating system over time, up to a predefined global allocation limit, and is then efficiently used by SAP HANA as needed.

SAP HANA preallocates and manages its own memory pool, used for storing in-memory table data, thread stacks, temporary results, and other system data structures. When more memory is required for table growth or temporary computations, the SAP HANA memory manager obtains it from the pool. When the pool can't 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.

The limit can be changed by modifying the `global_allocation_limit` configuration parameter in the global.ini file. It can be defined either as a fixed value in MB or as a flexible percentage of the available main memory size. If you enter a percentage value, the precise value of the limit is calculated automatically by the system. Moreover, if you then change the size of the container where the system runs the allocation limit automatically adjusts to the correct percentage of the new container size.

There's normally no reason to change the value of this parameter. However, on development systems with more than one SAP HANA system installed on a single host you could limit the size of the memory pool to avoid resource contentions or conflicts.

A change may also be necessary to remain in compliance with the memory allowance of your license if you purchased a license for less than the total amount of physical memory available. This change is illustrated in the following examples:

> **Example**
>
> - You have a server with 512 GB, 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).

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 can't 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). Therefore:
>
> - 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 allocation happens because service 1 is already using 50 GB and together they can't 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 only allocates memory for internal operations after 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, it's managed on a host-by-host basis; that is, column partitions are unloaded only on hosts with an acute memory shortage.

Avoid table (column or partition) unloading since it leads to performance degradation later when the table is queried and the data has to be reloaded. You can identify pool exhaustion by examining the `M_CS_UNLOADS` system view.

However, it's still possible that the memory manager needs more memory than is available leading to an out-of-memory failure. This may happen, for example, when too many concurrent transactions use up all memory, or when a complex query performs a cross join on large tables and creates a huge intermediate result that exceeds the available memory.

Related Information

Change the Global Memory Allocation Limit

9.1.16.1.4 SAP HANA Memory Usage and the Operating System

Due to the way in which SAP HANA manages memory, the relationship between Linux memory indicators and SAP HANA's own memory indicators may not correlate as expected.

From the perspective of the Linux operating system, SAP HANA is a collection of separate processes. Linux programs reserve memory for their use from the Linux operating system. The entire reserved memory

footprint of a program is referred to as its virtual memory. Each Linux process has its own virtual memory, which grows when the process requests more memory from the operating system, and shrinks when the process relinquishes unused memory. You can think of virtual memory size as the memory amount that the process has requested (or allocated) from the operating system, including reservations for its code, stack, data, and memory pools under program control. SAP HANA's virtual memory is logically shown in the following figure:

SAP HANA Virtual Memory

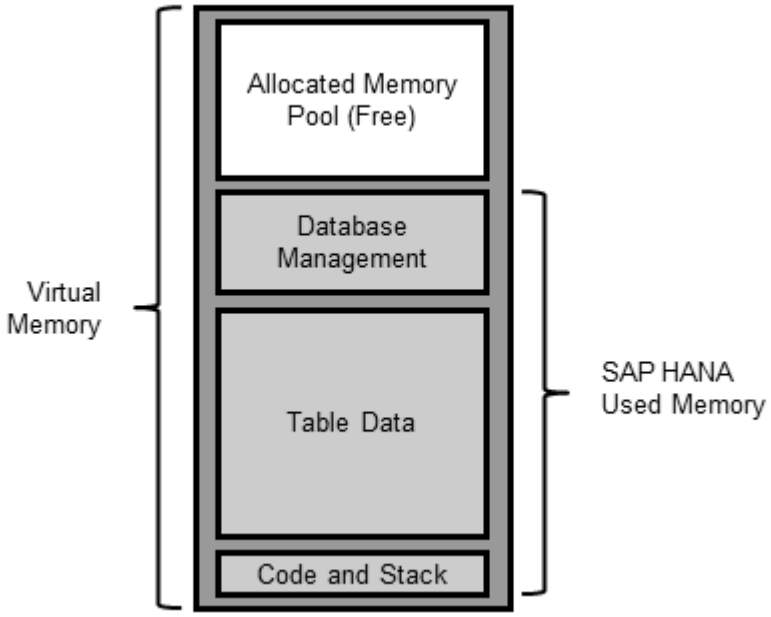

> **i** Note
>
> SAP HANA really consists of several separate processes, so the preceding figure shows all SAP HANA processes combined.

Virtual, Physical, and Resident Memory

When part of the virtually allocated memory actually needs to be used, it's 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's 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's enough physical memory so that swapping is disabled and isn't observed.

This memory usage can be illustrated as follows:

SAP HANA Resident Memory

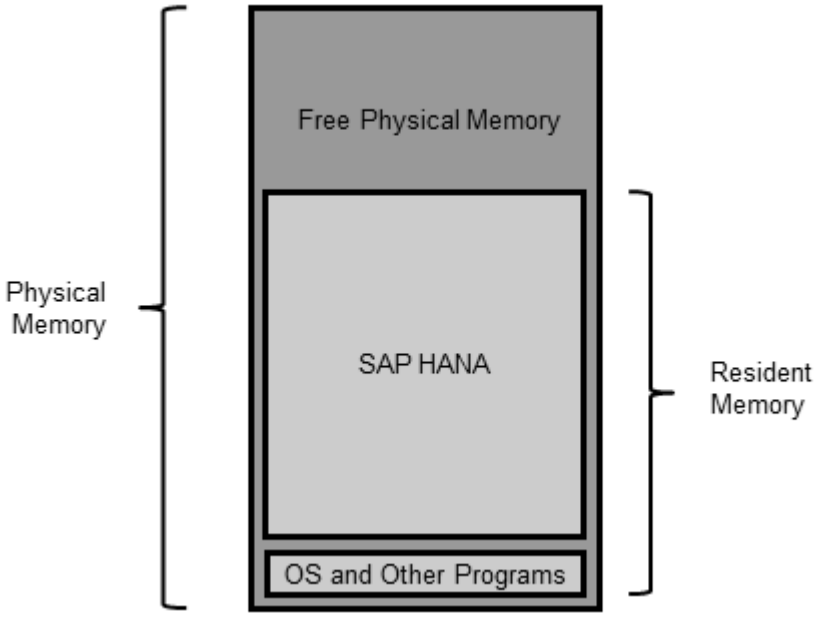

On a typical SAP HANA appliance, the resident memory part of the operating system and all other running programs usually doesn't exceed 2 GB. The rest of the memory is therefore dedicated for the use of SAP HANA.

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 can't satisfy the request, the SAP HANA memory manager requests and reserves 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 aren't affected. This creates a situation where the used memory value can shrink to below the size of SAP HANA's resident memory, which is normal.

> **i** Note
>
> The memory manager can also choose to return memory back to the operating system, for example when the pool is close to the allocation limit and contains large unused parts.

Related Information

SAP HANA Used Memory [page 406]
Memory Sizing [page 408]
Allocated Memory Pools and Allocation Limits [page 409]

9.1.16.2 Define CPU Cores Allocation Limits

As part of the provisioning process, you can ensure that CPU cores are shared appropriately between tenant databases. By setting the allocation limits, you ensure appropriate sharing of CPU cores, effectively specifying the maximum number of CPU cores that can be used by a particular tenant on a particular host.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

In the SAP HANA cockpit, you can modify the CPU cores allocation limit. This limit corresponds to the settings of the `max_concurrency` parameter in the global.ini file.

Each host in an SAP HANA system has a physical set of CPU cores. You can adjust the maximum number of CPU cores available for allocation at each of the following levels:

- System
- Tenant database (these settings override the system settings)
- Host (these settings override the database and system settings)

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Click *Global Allocation Limits* > *CPU Allocation*.

   For each database (or each database per host in a multi-host system), you can also refer to its mini-chart representing the daily CPU usage of the indexserver, where dark green shows the peak CPU usage, and light green shows the average CPU usage.
3. View the *Default CPU cores allocation limit*, or choose to edit the default by clicking the pencil icon. Changing the default causes each of the databases and hosts to inherit the value of the allocation limit (except those databases or hosts that have an allocation limit that already differs from the default).
4. View the *CPU Core Limit* for each of the databases, or choose to edit the value by clicking the pencil.
5. If you choose to edit the limit value, the *Edit Max CPU Cores* dialog displays. Enter a value, or clear the allocation limit input field to revert to the allocation limit inherited from the default value.

   The total number of CPU cores available for allocation is equal to the sum of the allocation limits for each database (regardless of whether you've edited the limit or allowed the default value to be inherited). A warning displays if this sum exceeds the number of physical CPU cores.

9.1.16.2.1 Applying NUMA Node Affinity Settings

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

You can use the `affinity` configuration parameter to restrict CPU usage of SAP HANA server processes to certain CPUs or ranges of CPUs.

Start by analyzing how the system CPUs are configured using the Linux `lscpu` command and then, based on the information returned, apply affinity settings in `daemon.ini` to bind specific processes to logical CPU cores. Processes must be restarted before the changes become effective. This approach applies primarily to the use cases of SAP HANA tenant databases and multiple SAP HANA instances on one server; you can use this approach, for example, to partition the CPU resources of the system by tenant database.

To make the changes described here, you require access to the operating system of the SAP HANA instance to run the `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 can't 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 following table gives a commentary on the most useful values based on an example system with 2 physical chips (sockets), each containing 8 physical cores. These cores are hyperthreaded to give a total of 32 logical cores.

#	Feature	Example Value
1	Architecture	x86_64
2	CPU op-modes	32-bit, 64-bit
3	Byte Order	LittleEndian

#	Feature	Example Value
4	CPUs	32
5	On-line CPUs list	0–31
6	Threads per core	2
7	Cores per socket	8
8	Sockets	2
9	NUMA nodes	2
21	NUMA node0 CPUs	0-7,16-23
22	NUMA node1 CPUs	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. Hyperthreading is where multiple threads are assigned to a single physical core In this example, there are 2 sockets, with each socket containing 8 physical cores (total 16). Two logical cores (#6) 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 nodes, one for each socket. Other systems may have multiple NUMA nodes per socket.
- 21–22: The 32 logical cores are numbered and assigned to one of the two NUMA nodes.

> **i Note**
>
> Even on a system with 32 logical cores and two sockets the assignment of logical cores to physical CPUs and sockets can be different. It's 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.

Command	Example Output	Commentary
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.

Other Linux commands that 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 can be used, for example, to isolate tenants) and `numactl` controls NUMA policy for processes or shared memory.

3. Based on the results returned you can use the `affinity` setting to restrict CPU usage of SAP HANA server processes to certain CPUs or ranges of CPUs. You can set restrictions for the following servers: nameserver, indexserver, compileserver, preprocessor, and xsengine (each server has a section in the daemon.ini file). The affinity setting is applied by the TrexDaemon when it starts the other HANA processes using the command `sched_setaffinity`. Changes to the affinity settings take effect only after restarting the HANA process. The following examples and commentary show the syntax for the ALTER SYSTEM CONFIGURATION commands required.

   Example

   To restrict the nameserver to two logical cores of the first CPU of socket 0 (see line 21 in the previous example), 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 previous example), use the following affinity settings:

   ```
   ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
   ('preprocessor', 'affinity') = '1-7,17-23'
   ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
   ('compileserver', 'affinity') = '1-7,17-23'
   ```

   Example

   To restrict the indexserver to all cores on socket 1 (see line 22 in the previous example), use the following affinity settings:

   ```
   ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini', 'SYSTEM') SET
   ('indexserver', 'affinity') = '8-15,24-31'
   ```

4. You can assign affinities to different tenants of a multi-tenant database on the same host as shown here. Run these SQL statements on the SYSTEMDB.

Example

In this scenario tenant NM1 already exists, here we add another tenant NM2:

```
CREATE DATABASE NM2 ADD AT LOCATION 'host:30040' SYSTEM USER PASSWORD
Manager1;
```

Set the configuration parameter to bind CPUs to specific NUMA nodes on each tenant. You can use the following notation with a dot to identify the specific tenant:

```
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini','SYSTEM')  SET
('indexserver.NM1', 'affinity') ='0-7,16-23';
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini','SYSTEM')  SET
('indexserver.NM2', 'affinity') ='8-15,24-31';
```

5. To assign affinities to multiple indexservers of the same tenant on the same host, execute the following SQL statements on the SYSTEMDB to apply the `instance_affinity[port]` configuration parameter:

Example

In this scenario, an indexserver is already running on tenant NM1 on port 30003. Here we add another indexserver on a different port:

```
ALTER DATABASE NM1 ADD 'indexserver' AT LOCATION 'host:30040';
```

Set the different instances of the `instance_affinity[port]` configuration parameter to bind CPUs to specific NUMA nodes on each indexserver. The configuration parameter has a 1–2 digit suffix to identify the final significant digits of the port number, in this example 3000**3** and 300**40**:

```
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini','SYSTEM')  SET
('indexserver.NM1', 'instance_affinity[3]')='0-7,16-23';
ALTER SYSTEM ALTER CONFIGURATION ('daemon.ini','SYSTEM')  SET
('indexserver.NM1', 'instance_affinity[40]')='8-15,24-31';
```

Restart the indexserver processes to make the affinity settings effective.

6. You can test the settings either in SQL or using hdbcons as shown here:

   Run this query on the tenant or SystemDB:

   ```
   select * from M_NUMA_NODES;
   ```

   Using hdbcons the process ID of the indexserver process is required as a parameter:

   ```
   hdbcons -p <PID> "jexec info"
   ```

Related Information

Memory and CPU Usage for Tenant Databases [page 403]
SAP HANA Monitoring Views for CPU Topology Details
SQL Statements to Apply NUMA Location Preferences
SAP Note 2470289

9.1.17 Monitoring Tenant Databases in SAP HANA Cockpit

As the tenant database administrator, you can monitor the availability, resource usage, and performance of tenant databases in the SAP HANA cockpit from the system database.

Aggregate database information is available on the *Database Overview* page of the system database. The *Database Overview* page of the system database has separate sections for *Tenant Monitoring and Administration* and *System DB Monitoring and Administration*. The latter displays information and links for the monitoring the system database as a resource itself. For more information, see *Using the System Overview*.

Drilling down on the *Overall Tenant Statuses* card displays the *Manage Databases* page, which provides you with further information for all tenant databases. From *Manage Databases*, you can then drill down to more detailed information about each individual tenant database. You can use the cockpit to monitor and manage more than one database, each running version SAP HANA 1.0 SPS 12 or later. Any database running version SAP HANA 2.0 SPS 01 or later is set in multiple-container mode by default. The cockpit can also monitor single-container systems running earlier versions of SAP HANA. When you access the *Manage Services* application, the operations you have the option to perform depend on whether you are displaying a tenant or a system database.

> **i** Note
>
> To perform operations on a tenant database, you have the system privilege DATABASE ADMIN.

Related Information

Using the Database Overview Page to Manage a Database [page 109]
Database Details [page 418]
Monitor Alerts for a Tenant Database [page 421]

9.1.17.1 Database Details

The *Manage Databases* page provides you with detailed information about all databases, as well as several drill-down options for more detailed information about individual databases.

The following table lists the information available for databases, as well as the available drill-down option.

Column	Description	Drill-Down Option
Database Name	Name of the tenant database	Click the database name to open the Overview, from which you can drill down to the *Manage Services* page. *Manage Services* allows you to analyze the status and resource usage of the individual services of the database. For more information, see *Service Details*. From *Manage Services*, you can also stop and start services.
Status	Status of the database: • Redistribution Running: A table redistribution plan has been executed • Deleting: Database is being deleted • Error: Database is in an error state • Running with Issues: At least one service isn't running, or there is at least one high alert • Maintenance: Database is in enforced maintenance mode • Replicating: Database is undergoing replication • Running: Database is running • Starting: Database is starting • Stopped: Database is stopped • Stopping: Database is stopping • Transitioning: Database is transitioning (starting or stopping) • Unknown: Database is in an unknown state	No specific drill-down
Start Time	The time of the most recent start of the database	No specific drill-down

Column	Description	Drill-Down Option
Alerts	The number of high and medium priority alerts in the database	Click the number of alerts to open them on the *Alerts* page *Alerts* allows you to view and analyze alerts occurring in the database. You can view past occurrences of alerts. For more information, see *Alert Details* and *Alert Priorities*. You can also see how alerts are configured. For more information, see *Alert Checker Details* and *Alert Checker Statuses*. **i Note** Only those alerts that identify situations with a potentially system-wide impact are visible, for example, the physical memory on a host is running out. Alerts that expose data in the tenant database (for example, table names) are **not** visible to the system administrator in the system database.
Used Memory	The used memory of the database in relation to the system	Click the used memory bar to open the *Performance Monitor* app The *Performance Monitor* app allows you to visually analyze historical performance in the database across a range of related performance indicators. For more information, see *Key Performance Indicators*.
CPU Usage	The CPU usage of the database in relation to the system	Click the CPU usage bar to open the *Performance Monitor* app The *Performance Monitor* app allows you to visually analyze historical performance in the database across a range of related performance indicators. For more information, see *Key Performance Indicators*.
Disk Usage	The disk usage of the database in relation to the system	Click the disk usage bar to open the *Performance Monitor* app The *Performance Monitor* app allows you to visually analyze historical performance in the database across a range of related performance indicators. For more information, see *Key Performance Indicators*.
Fallback Snapshot	Lists whether a fallback snapshot is available. A fallback snapshot can be useful if you perform changes to the contents of a database that you need to roll back quickly.	No specific drill-down

Related Information

Service Details [page 160]
Alert Details
Alert Priorities
Alert Checker Details
Alert Checker Statuses
Key Performance Indicators [page 519]

9.1.17.2 Monitor Alerts for a Tenant Database

Alert situations in tenant databases can potentially impact the health of the overall system. For this reason, you as system administrator can monitor alerts occurring in individual tenant databases. You can monitor from the system database in the SAP HANA cockpit.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- To be able to drill down to the alert information of a tenant database, you must have registered it as a database in the cockpit.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. Open the alerts of a particular database by clicking the number of alerts indicated for that database in the *Alerts* column.

   > → Tip
   >
   > You can also access database-specific alerts from *Alerts* in the *Manage Databases* app.

Next Steps

It can be helpful to see how alerts are configured in individual tenant databases. To navigate to the configuration of alert checkers from the *Alerts* app, click *Alert Definitions* on the top right-hand corner.

Related Information

Using the Database Overview Page to Manage a Database [page 109]
Services [page 159]
Work with Alerts

9.1.18 Add or Remove Services in a Tenant Database

Add or remove services in the tenant database.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Add a Service

Prerequisites

- You are be connected to the system database.
- You have:
  - The system privilege DATABASE ADMIN.
  - The EXECUTE privilege on the stored procedure SYS. UPDATE_LANDSCAPE_CONFIGURATION.

Context

You can only add services when connected to the SYSTEMDB.

Procedure

1. On the *Database Overview* of the system database, click *Databases Management*.
2. In the *Status* column, click the status message of the tenant being managed.

   The *Manage Services* page for the tenant opens.

> ⚠ Caution
>
> If you click the database name instead of the status message, you are disconnected from SYSTEMDB and connected to the tenant. The *Database Overview* page for the tenant opens instead of the *Manage Services* page.

3. Click *Manage Services* > *Add Service*.
4. Select a service from the list.
5. Select a host and enter its port, or allow a host to be auto-assigned.

Results

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.

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

There are other services available to add or remove if you've installed an SAP HANA capability (such as SAP dynamic tiering), and the host has been added to the SAP HANA SYSTEM DB.

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.

Remove a Service

Prerequisites

- You have the RESOURCE ADMIN and CATALOG READ system privileges.
- You have the SELECT privilege on _SYS_STATISTICS

Context

You can remove a service from a tenant when connected to either the SYSTEMDB or the tenant. You access the *Manage Services* page for the tenant differently depending on your connection method.

Procedure

1. Choose one of the following options:

Method	Action
Connected to SYSTEMDB	1. At the top of the *Database Overview* page, click *Database Management*. 2. In the *Status* column, click the status message of the tenant being managed.
Connected to the tenant	1. On the *Database Overview* with the *Monitoring* or *All* view selected, click *Manage Services* on the *Services* card.

2. Select the service to remove.

   Not all services can be removed. You can't remove a global service, the master indexserver, or the primary indexserver on a host.

3. Choose one of the following options:

Method	Action
Connected to SYSTEMDB	Click *Manage Services* > *Remove Service*.
Connected to the tenant	Click *Remove Service*.

4. Confirm the removal.

Results

- Data volumes and traces files are removed. The data of the index server is distributed across the remaining index server instances.
- The service is removed from the M_SERVICES system view.
- The service is stopped and removed from the system landscape information of system database.

Related Information

ALTER DATABASE Statement (Tenant Database Management)

9.1.19 Start and Stop Services in a Tenant Database

Start, stop, or kill services in the tenant database.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Start Missing Services

Prerequisites

- You are be connected to the system database.
- You have the system privilege DATABASE ADMIN.

Context

You can only start missing services when connected to the SYSTEMDB.

Procedure

1. On the *Database Overview* of the system database, click *Databases Management*.
2. In the *Status* column, click the status message of the tenant being managed

   The *Manage Services* page for the tenant opens.

   > ⚠ Caution
   >
   > If you click the database name instead of the status message, you are disconnected from SYSTEMDB and connected to the tenant. The *Database Overview* page for the tenant opens instead of the *Manage Services* page.

3. Click *Stat Missing Services*.

Stop or Kill a Service

Prerequisites

- You have the system privilege RESOURCE ADMIN.

Context

You can stop or kill a service from a tenant when connected to either the SYSTEMDB or the tenant. You access the *Manage Services* page for the tenant differently depending on your connection method.

Procedure

1. Choose one of the following options:

Method	Action
Connected to SYSTEMDB	1. At the top of the *Database Overview* page, click *Database Management*. 2. In the *Status* column, click the status message of the tenant being managed.
Connected to the tenant	1. On the *Database Overview* with the *Monitoring* or *All* view selected, click *Manage Services* on the *Services* card.

2. Select the service to stop or kill.
3. Click *Kill Service* or *Stop Service*.
4. Confirm the stop or kill.

9.1.20 Change the Port of a Service in a Tenant Database

You can change the port of a service in the SAP HANA cockpit.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- The database user (with which you've connected to the database) must have the system privilege DATABASE ADMIN.

Procedure

1. At the top of the *Database Overview* page, click *Database Management*.
2. To change the port of a service, select *Manage Services* > *Change Port*.
3. In the *Change Port* dialog, select the *New Port* for the service.
4. Optional: If reserved instances are configured for the host, you can choose to select ports available on another instance.

   > **i** Note
   >
   > The default port number range for tenant databases is 3<instance>40—3<instance>99. This means that the maximum number of tenant databases that can be created per instance is 20. However, you can increase this by reserving the port numbers of further instances.

5. Optional: Start the service after changing the port.

   By default, the service isn't started after the port is changed. If you want the service to be started after the port is changed, enable the *Start Automatically* option.

Related Information

ALTER DATABASE Statement (Tenant Database Management)

9.2 Database Administration

Administer your system or tenant database.

9.2.1 Configure Host Failover

For multi-host systems, you can configure host auto-failover so that if an active host fails, standby hosts take over to ensure the continued availability of the database.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Host roles for failover are normally configured during installation. Using SAP HANA cockpit, you can monitor the status of individual hosts and switch the configured roles of hosts; you can't increase or decrease the number of worker hosts and standby hosts in relation to each other.

The primary reason for changing the configured roles is to prepare for the removal of a host. In this case, change the configured role of the name server host to SLAVE and the configured role of the index server host to STANDBY before stopping the database instance on the host and removing the host.

> **i** Note
>
> To change host configuration, your database user must have the system privilege RESOURCE ADMIN and the object privilege EXECUTE on the procedure UPDATE_LANDSCAPE_CONFIGURATION.

Procedure

1. On the *Database Overview*, with the *Administration* or *All* view selected, click the *Host Failover* link on the *Database Administration* card.
   All the hosts in the system are displayed, whether or not they're operational, as well as additional information about their auto-failover status and configuration.
2. Click the gear button to customize which columns to display.

Column	Description
Host	Host name
Active	Indicates the status of services running on the host. The following statuses are possible: ○ YES All services are active. ○ PARTIAL Some services are active. ○ STARTING Some services are active, and some are starting. ○ STOPPING Some services are active, and some are stopping. ○ NO No services are active.

Column	Description
Host Status	Indicates the host's status and whether the system is operational. The following statuses are possible: ○ OK The system is operational and the host's actual role corresponds to its configured role. ○ IGNORE The system is operational. The host is configured as a standby host and is available, but not in use. ○ INFO The system is operational. The host's actual role is different from its configured role. ○ WARNING The system isn't operational. The host will become available after start-up or failover. ○ ERROR The system isn't operational. The host is missing.
Failover Status	Displays the failover status so you can see which hosts are active and which are on standby. The following statuses are possible: ○ <Empty> Failover isn't active or pending. ○ WAITING ... SEC The host has failed. The system is waiting to fail over. ○ WAITING The host has failed. The system is waiting for the host to restart to prevent unnecessary failover. ○ FAILOVER TO `<host>` The host has failed and failover to a target host is in progress. ○ FAILBACK TO `<host>` Failback to a worker host is in progress. Failback happens when the assigned standby host is stopped. However, there's no automatic failback while the standby host is still assigned since it would cause downtime. ○ FAILED Failover isn't possible, for example, no further standby hosts available. For more information, see the nameserver trace.
Nameserver Role (Configured)	Specifies the host's configured role as name server. The following roles are possible: ○ MASTER 1, MASTER 2, MASTER 3 When you install a distributed system, up to three hosts are automatically configured as master name servers. The configured nameserver role of these hosts is MASTER 1, MASTER 2, and MASTER 3. ○ SLAVE Additional hosts in your system are configured as slave name servers. The configured nameserver role of these hosts is SLAVE.

Column	Description
Nameserver Role (Actual)	Specifies the host's actual role as name server. The following roles are possible: ○ MASTER During system start-up, one of the hosts configured as master name servers (that is, those hosts with the configured name server role MASTER 1, MASTER 2, or MASTER 3) is designated to be the active master name server. The actual nameserver role of this host is MASTER. This master name server assigns one volume to each starting index server (with actual role MASTER or SLAVE), or no volume if it's a standby host (actual indexserver role STANDBY). If this active master nameserver host fails, one of the remaining hosts configured as a master nameserver becomes the active master name server. ○ SLAVE The actual nameserver role of the remaining hosts configured as master and slave hosts is SLAVE.
Indexserver Role (Configured)	Specifies the host's configured role as index server. The following roles are possible: ○ WORKER ○ STANDBY 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 isn't used for database processing. All database processes run on the standby host, but they're idle and don't allow SQL connections.

Column	Description
Indexserver Role (Actual)	Specifies the host's actual role as index server. The following roles are possible: ○ MASTER The actual master indexserver is assigned on the same host as the name server with the actual role MASTER. The actual index server role of this host is MASTER. The master index server provides metadata for the other active index servers (with actual indexserver role SLAVE). ○ SLAVE The actual index server role of remaining hosts (except hosts configured as standby hosts) is SLAVE. These hosts are active index servers and are assigned to one volume. If an active index server fails, the active master name server assigns its volume to one of the standby hosts. ○ STANDBY The actual indexserver role of standby hosts is STANDBY. A standby host isn't assigned a volume by the active master name server and it doesn't open an SQL port. 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. **i Note** Failover is configured only for the name server and the index server on each host. The other components (for example, xsengine) aren't configured individually as they're always failed over together with the index server.
Failover Group (Configured/Actual)	A failover group can be defined for each host. If there's a failover, the name server tries to fail over to a host within the same group.
Worker Groups (Configured/Actual)	The worker groups (also referred to as host subroles) for the host can be set here. The groups are required to support heterogeneous hardware in the landscape that is required, for example, for the extension node feature. Worker groups may also be relevant in a single-host installation. The worker group name is a free text value that is validated to trap illegal characters.

Column	Description
Host Roles (Configured)	Specifies the host's configured database role. The following roles are possible: ○ WORKER Worker host for database processing ○ STANDBY Standby host for database processing Depending on your installation, the following additional host roles can be configured: ○ EXTENDED_STORAGE_WORKER Worker host for SAP HANA dynamic tiering ○ EXTENDED_STORAGE_STANDBY Standby host for SAP HANA dynamic tiering ○ ETS_WORKER Worker host for SAP HANA accelerator for SAP ASE ○ ETS_STANDBY Standby host for SAP HANA accelerator for SAP ASE ○ STREAMING Host for SAP HANA Streaming Analytics ○ XS_STANDBY Standby host for SAP HANA XS advanced runtime ○ XS_WORKER Host for SAP HANA XS advanced runtime **i Note** Multiple host roles aren't supported in production environments. However, if XS advanced runtime is installed, hosts can share multiple roles.
Host Roles (Actual)	Specifies the host's actual database role
Storage Partition	Specifies the number of the `mnt000...` subdirectory used by the host for storing data and logs, for example, *1* if the subdirectory is `mnt00001`, *2* if it's `mnt00002`, and so on. During installation, volumes for storing data and log files are defined, which are the directories where data and logs are stored. The default directories are: ○ `/hana/data/<SID>` for data ○ `/hana/log/<SID>` for logs Each active host has exactly one subdirectory beneath these directories called `mnt00001`, `mnt00002`, and so on. The next level in the file hierarchy is the actual volume, with one subdirectory for each service called `hdb00001`, `hdb00002`, and so on. In the event of failover, the volumes of the failed host are reassigned to the standby host.

Column	Description
Removal Status	Indicates the status of the table redistribution operation used to move data off the index server of a host that you plan to remove. Before you can remove an active host from a single-container system, you must move the tables on the index server of this host to the index servers on the remaining hosts in the system. Once the value in the *Removal Status* column changes to REORG FINISHED or REORG NOT REQUIRED, you can physically remove the host using the SAP HANA lifecycle management tool hdblcm(gui). If your system is configured as a multiple-container system, you have to remove tenant-specific services first and then remove the host using the SAP HANA database lifecycle manager (HDBLCM). For more information, see *Remove a Service from a Tenant Database* in the *SAP HANA Administration Guide*. The following statuses are possible: ○ <Empty> Host hasn't been marked for removal. ○ REORG PENDING A redistribution operation is required to move tables to other hosts. ○ REORG ACTIVE A redistribution operation is in progress. For more information, you can query the system tables SYS.REORG_OVERVIEW and SYS.REORG_STEPS. ○ REORG FAILED A redistribution operation was executed and failed. For more information, query the system table SYS.REORG_STEPS. ○ REORG FINISHED A redistribution operation has completed. The host can be uninstalled. ○ REORG NOT REQUIRED A redistribution operation isn't required. The host can be uninstalled.

3. If you change a configured role or configured group, click *Apply* so that your changes take effect.

Related Information

Multiple-Host (Distributed) Systems
Add or Remove Services in a Tenant Database [page 422]
Monitoring Tenant Databases in SAP HANA Cockpit [page 418]

9.2.2 License Management

A valid license key is required to use the SAP HANA database. Additional license keys are required for certain applications running on SAP HANA, as well as and certain SAP HANA options and capabilities. You can use the SAP HANA cockpit to see which licenses are available in your system, to install new license keys, and to view memory usage with respect to licensing.

Related Information

Important Disclaimer for Features in SAP HANA
View Licenses [page 436]
Install a Permanent License [page 439]
Delete Licenses [page 442]
License Keys for SAP HANA Database [page 434]

9.2.2.1 License Keys for SAP HANA Database

The SAP HANA database supports two kinds of license keys: temporary license keys and permanent license keys

SAP HANA licenses can be installed for the system database (global) or for a single tenant database (local). Global licenses are for the system database and all the tenant databases, but a license installed in a tenant database governs only that tenant database. If you remove a tenant-specific license key, that tenant database reverts to the global license key installed in the system database.

Temporary License Keys

A temporary license key is automatically installed with a new SAP HANA system. A temporary license key is valid for 90 days. During this period, request and install a permanent license key.

Permanent License Keys

Permanent License Keys: Expiration

Before a permanent license key expires, request and apply a new permanent license key. If a permanent license key expires, a (second) temporary license key is automatically installed. This temporary license key is valid for 28 days. During this time, you can request and install a new permanent license key.

You can request a permanent license key on SAP Support Portal (http://support.sap.com) under *Request Keys*. Permanent license keys are valid until the predefined expiration date. Furthermore, they specify the amount of memory licensed to an SAP HANA installation.

Permanent License Keys: Types

There are two types of permanent license key for SAP HANA: unenforced and enforced.

If an unenforced license key is installed, the operation of SAP HANA isn't 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 lockdown occurs, either SAP HANA needs to be restarted, or a new license key that covers the amount of memory in use needs to be installed.

The two types of permanent license key differ from each other in the following line in the license key file:

License Key Type	License Key File Entry
Unenforced	SWPRODUCTNAME=SAP-HANA
Enforced	SWPRODUCTNAME=SAP-HANA-ENF
	SWPRODUCTNAME=SAP-HANA-DEV
	SWPRODUCTNAME=SAP-HANA-DIGITAL

> **i** Note
>
> It's technically possible to install an enforced 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, excessive memory consumption doesn't result in a system lockdown. However, if one license key expires and becomes invalid, the other license, if valid, becomes the valid license key of the instance. If the latter is an enforced license key, then the memory consumption is checked.

License Keys for Tenant Databases

You can install permanent license keys in individual tenant databases. The license key installed in a tenant database is valid for that database only and takes precedence over the license key installed in the system database. If a tenant-specific license key isn't installed, the system database license key is effective in the tenant database.

> **→** Tip
>
> The system view SYS.M_LICENSE provides tenant administrators with information on the license key effective in their tenant database, as well as where the license key is installed: in the tenant database itself or in the system database. System administrators can use the view SYS_DATABASES.M_LICENSE to see the same information for all tenant databases.

System Lockdown

The system goes into lockdown mode in the following situations:

- The permanent license key has expired and either:
  - You didn't renew the temporary license key within 28 days, or
  - You did renew the temporary license key but the hardware key has changed
- The installed license key is an enforced license key and the current memory consumption exceeds the licensed amount plus the tolerance.
- You deleted all license keys installed in your database.

In lockdown mode, it isn't possible to query the database. Only a user with the system privilege LICENSE ADMIN can connect to the database and execute license-related queries, such as, obtain previous license data, install a new license key, and delete installed license keys.

In addition, the database can't be backed up in lockdown mode.

Additional SAP HANA Licenses

Additional licenses are required for certain applications running on the SAP HANA database, as well as certain SAP HANA options and capabilities. For more information, see SAP Note 1644792 (License key/installation of SAP HANA).

Related Information

SAP Support Portal
SAP Note 1644792

9.2.2.2 View Licenses

With SAP HANA cockpit, you can view licenses installed in your SAP HANA database.

Prerequisites

- You have the system privilege LICENSE ADMIN.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, navigate to the the *Database Administration* card, then choose *Database Licenses*.

   All licenses installed in the database are listed.
2. Choose an option:

Option	Description
To view more information about a license	Select the license. License-specific information, and, if applicable, memory usage data, is displayed. For more information, see *License Details*.

Option	Description
	Note If you're viewing license information in the system database, memory usage data is for the system as a whole. In a tenant database, only the usage data of the tenant database itself is shown.
To view information about the licenses for one or more tenant databases	By default, the licenses for the system database are shown. To view a different database, select it from the *Database Name* menu. To view licenses for the system database and all tenant-specific licenses, remove all database names from the *Database Name* field. Here, you can see whether a tenant database is using the system database license, or whether it has its own license. If a license is uninstalled, it is no longer visible.
To request a new license	Select the database, then choose *Request New License*. The following information is displayed: Hardware key, System ID, installation number, and system number. This information is needed to apply for a new license. To apply for an SAP HANA license, choose *Go to the SAP Support Portal*. For more information, see *SAP Note 1644792 (License key/installation of SAP HANA)*.

Related Information

License Details [page 438]
License Keys for SAP HANA Database [page 434]
SAP Note 1644792

9.2.2.2.1 License Details

The *Database Licenses* page provides you with detailed information about all licenses installed in the SAP HANA database.

General Information for SAP HANA Database Licenses

Field	Description
Hardware Key	Unique hardware key
System ID	Unique SAP system identifier
License Type	License type, either permanent or temporary
Product Description	SAP HANA database
Starts On	Date as of which the license is valid
Expires On	Date on which the license expires
Licensed Memory Usage	Amount of memory usage licensed
Peak Memory Usage	Highest recorded value for main memory usage consumed by the SAP HANA database
GLAS ID	Unique product ID of SAP HANA required for license auditing

> **i Note**
>
> If you're viewing the information in the system database, memory usage data is for the system as a whole. In a tenant database, only the memory usage data of the tenant is shown.

The following additional fields are available if SAP Business Warehouse (SAP BW) is running on the SAP HANA database.

Field	Description
Peak Memory Usage of SAP BW	Highest recorded value for main memory usage consumed by SAP BW
GLAS ID of SAP BW	Unique product ID of SAP BW required for license auditing
Peak Memory Usage of Non-SAP BW Components	Highest recorded value for main memory usage consumed by SAP HANA
GLAS ID of Non-SAP BW Components	Unique product ID of SAP HANA required for license auditing

General Information for Other SAP HANA Licenses

Field	Description
Hardware Key	Unique hardware key
System ID	Unique SAP system identifier
License Type	License type, either permanent or temporary
Product Description	Product for which the license is valid
Starts On	Date as of which the license is valid
Expires On	Date on which the license expires
Licensed `<license_metric>`	Amount of usage licensed The unit of measurement varies from product to product. For example, SAP HANA dynamic tiering licensing is based on hard disk usage.
Peak `<license_metric>`	Highest recorded usage value consumed
GLAS ID	Unique product ID required for license auditing

Memory Usage

The memory usage graph is available for SAP HANA database licenses and shows you the peak memory usage recorded every month for the previous twelve months.

9.2.2.3 Install a Permanent License

To use SAP HANA, you must request and install a permanent license key. You can do this in the SAP HANA cockpit.

Prerequisites

- You have authorization to request permanent license keys on SAP Support Portal.
- You have the system privilege LICENSE ADMIN.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Request and install a permanent license key, for example, if the current license key of your SAP HANA system is about to expire, or you want to extend the amount of memory licensed for your system.

SAP HANA licenses can be installed for:

- The system database (global)
  Global licenses are for the system database and all the tenant databases.
- A single tenant database (local)
  A license installed in a tenant database governs only that tenant database.
  If you remove a tenant-specific license key, that tenant database reverts to the global license key installed in the system database.

> **i Note**
>
> You can also request and install the license keys required for applications running on SAP HANA, as well as SAP HANA options and capabilities, using the procedure described here.

Procedure

1. From the *Database Overview*, navigate to the *Database Administration* card, then choose *Database Licenses*.

   > **i Note**
   >
   > If the system is in lockdown, you're automatically prompted to navigate to the *License Details* page.

2. Request a permanent license key as follows:
   a. On the *Database Licenses* page, choose *Request New License*.

      If the system is currently running on a temporary license key, the hardware key and the system ID are displayed.

      If the system already has a valid permanent license key, the installation number and system number are displayed. You need this information to fill out the request form for the license key.
   b. Choose *Go to SAP Support Portal*.
   c. In SAP Support Portal, choose *Request Keys*.

      You're forwarded to the license key application of the SAP ONE Support Launchpad, where you can request a new key. When completing the request form, if you have the installation number and system number, then enter them first so that the other input fields are auto-completed. When you've finished, choose *Submit*.

      The permanent license is sent to you as an e-mail attachment.
3. To install the license key, choose *Upload System DB License*, then upload the license file (`*.txt` file) that you received by e-mail.

   > **i Note**
   >
   > If you're 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.

4. If required, specify the usage type for the system.

   If a usage type has already been defined, you aren't prompted to specify a usage type.

   The usage type can be:

   - NONPRODUCTION
     Choose this type if your system isn't used for production.
     For example, for a system used for testing or quality assurance.
   - RUNTIME
     Choose this type if you're using a RUNTIME license from SAP.
   - FULL-USE
     Choose this type if you're using a FULL-USE license from SAP.
   - ADMINISTRATION
     Choose this type if you're using your system as the basis to run SAP HANA Cockpit.

   Optionally, use the following SQL statement to set the usage type when installing the license key:

   ```
   SET SYSTEM LICENSE <license_key> [USAGE [NONPRODUCTION|RUNTIME|FULL|
   ADMINISTRATION]]
   ```

9.2.2.4 Export Usage Data

SAP HANA licensing is based on the amount of memory used. The SAP HANA database tracks the highest value for memory used per calendar month for a one-year period. You may be requested by SAP to export this data for license audit purposes to verify the suitability of your license.

Prerequisites

- You have the system privilege LICENSE ADMIN.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. From the *Database Overview*, navigate to the *Database Administration* card, then choose *Database Licenses*.
2. Choose *Export System DB System Measurement*.
3. If a usage type is not already defined, specify a usage type.

   For more information, see *Install a Permanent License*.

Results

Usage data for all installed licenses is exported to the file `SAPHANASystemMeasurement.xml`, which is downloaded in line with your browser's file download settings.

> **i** Note
>
> In Safari, the file isn't automatically downloaded. Instead, the content opens in a new tab or window and you must manually save the file by pressing `CMD` + `S`, choosing *page source*, and specifying a file name.

Related Information

SAP Note 1704499
Install a Permanent License [page 439]

9.2.2.5 Delete Licenses

You can delete all existing license keys in the SAP HANA database, or you can delete licenses in the system database, leaving only tenant-specific licenses.

Prerequisites

- You have the system privilege LICENSE ADMIN.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, navigate to the the *Database Administration* card, then choose *Database Licenses*.
2. Choose the option to delete all licenses.

If you are Connected To...	Option
A tenant database	*Delete All Licenses* The license keys installed in the system database now take effect.

If you are Connected To...	Option
The system database	*Delete All System DB System Licenses* The system database is locked down. Any tenant databases that do not have their own tenant-specific license key are also locked down. To unlock the databases, a new permanent license key is required.

Results

All permanent license keys are deleted.

9.2.3 Database Configuration in SAP HANA Cockpit

From the *Database Overview*, you can drill down to view and manage configuration (`*.ini`) files.

To open the *Database Configuration* page, on the *Database Overview* page, with the *Administration* or *All* view selected, click the *Manage Database Configuration* link on the *Database Administration* card.

Configuration files are separated into sections; sections bundle parameters of the same category. Parameters can be configured at different levels or layers depending on the configuration file. Layers control where parameter values apply and how parameters inherit default values. The following layers are available:

Layer	Description
Default	The default value for the parameter
System	The system-specific value for the parameter (configurable in the system database) If a system-specific value isn't configured for a parameter, then the default value applies.
Database	The database-specific value for the parameter (configurable in the system or tenant database) For some parameters, it's possible to set database-specific values. If a database-specific value isn't configured, then the host-specific value applies.
Host	The host-specific value for the parameter (configurable in the system database) For some parameters, it's possible to set host-specific values for multiple-host systems. If a host-specific value isn't configured for a parameter that can be set at host level, the system-specific value applies.

You can change configuration values (*Change Layer*) or quickly modify a value or assign a default.

You can filter by *Configuration File*, *Section*, and *Host* in order to display specific configuration file contents.

In general, SAP recommends that you don't change the default values of parameters unless the documentation suggests it or you're instructed to do so by SAP Support. While most parameters can be changed when the

database is running, changes to some parameters require a database restart to take effect. To find out whether a restart is required for frequently used parameters, refer to the online reference in the SAP Help Portal.

Related Information

Database-Specific Configuration Parameters [page 444]
Add a System Property Section [page 447]
Add a System Property Parameter [page 448]
Modify a System Property in SAP HANA Cockpit [page 449]
Restore a System Property Default in SAP HANA Cockpit [page 450]
SAP HANA Configuration Parameter Reference

9.2.3.1 System Properties

You can add new sections and parameters to configuration files, at one or more layers. You can also override the default value for existing properties.

9.2.3.1.1 Database-Specific Configuration Parameters

In addition to the layers "default", "system", and "host", system configuration files also have a "database" layer to facilitate the configuration of properties for individual databases.

In general, you can configure database-specific properties both in the system database and in tenant databases themselves. Properties configured in the system database can be applied to all databases (if configured in the system layer) or to specific databases (if configured in database layer).

Properties configured in a tenant database apply to that tenant database only. Only properties in the following files can be configured in tenant databases:

- `attributes.ini`
- `docstore.ini`
- `dpserver.ini`
- `esserver.ini`
- `executor.ini`
- `extensions.ini`
- `global.ini`
- `indexserver.ini`
- `multidb.ini`
- `scriptserver.ini`
- `xsengine.ini`

File Location

If properties are configured in the database layer, a database-specific configuration file is stored at the following location on the server: `/hana/shared/$SID/global/hdb/custom/config/DB_<dbname>`

> **Example**
>
> The properties in the `nameserver.ini` file aren't 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're connected to a tenant database, you see that the database-layer value of a property is also displayed as the system-layer value. This value occurs because, from the perspective of the tenant database, the database and the system are effectively the same. The true system-layer value (that is, the value configured for all databases in the system database) is displayed in the tenant database as the default-layer value.

Values configured in the host layer take precedence over database-layer values. Host values can only be configured in the system database.

To view actual configuration values, on the resources *Database Overview* page, select the *Manage System Configuration* link on the *Database Administration* card, or query the following system views:

- M_INIFILE_CONTENTS (SYS_DATABASES)
  This view can be accessed only from the system database. It contains the values configured for all properties on system, host, and database layer for **all** active databases.
- M_INIFILE_CONTENTS (SYS)
  This view is available in every database and contains the values that apply to the database in question. Values that were configured in the system layer in the system database are identified as default-layer values. Values that were configured in the database layer in the tenant database are identified as system- and database-layer values. Values configured at the host layer are shown only for hosts on which the database is running.

Example

A system has 3 tenant databases DB1, DB2, and DB3, distributed across 2 hosts Host A and Host B:

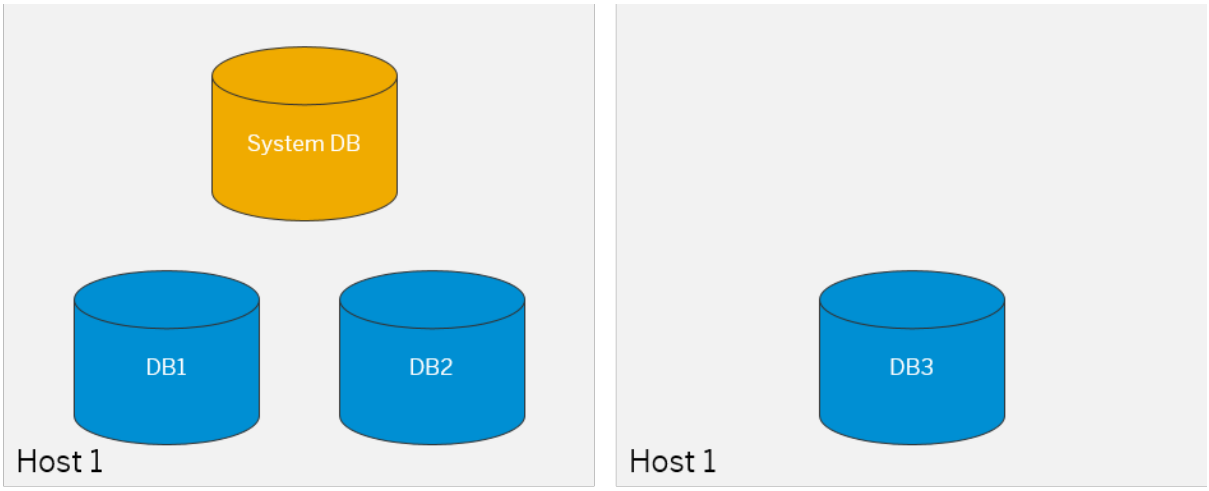

The default value of the property `[execution] max_concurrency` in the `global.ini` file is `0`. The system administrator changes the default configuration of this property in the `indexserver.ini file` as follows:

First, the system administrator creates a new system-layer value (**10**) in `indexserver.ini`. Since the system-layer value applies to all tenant databases and can't be changed by a tenant database user, users on all tenant databases initially see the value `10` as the default configuration:

Executed in DB1, DB2, and DB3:

```
SELECT * FROM SYS.M_INIFILE_CONTENTS WHERE KEY='max_concurrency'

| FILE_NAME        | LAYER_NAME   | TENANT_NAME   | HOST       | SECTION    | KEY             | VALUE |
|------------------|--------------|---------------|------------|------------|-----------------|-------|
| global.ini       | DEFAULT      |               |            | execution  | max_concurrency | 0     |
| indexserver.ini  | DEFAULT      |               |            | execution  | max_concurrency | 10 <--- Effective value
```

Next, the system administrator sets a new value (**20**) for DB1, while leaving the configuration for DB2 and DB3 unchanged.

Executed in DB1:

```
SELECT * FROM SYS.M_INIFILE_CONTENTS WHERE KEY='max_concurrency'

| FILE_NAME        | LAYER_NAME   | TENANT_NAME   | HOST       | SECTION    | KEY             | VALUE |
|------------------|--------------|---------------|------------|------------|-----------------|-------|
| global.ini       | DEFAULT      |               |            | execution  | max_concurrency | 0     |
| indexserver.ini  | DEFAULT      |               |            | execution  | max_concurrency | 10    |
| indexserver.ini  | SYSTEM       |               |            | execution  | max_concurrency | 20 <-\
| indexserver.ini  | DATABASE     |               |            | execution  | max_concurrency | 20 <--- Effective value
```

Executed in DB2 and DB3:

```
SELECT * FROM SYS.M_INIFILE_CONTENTS WHERE KEY='max_concurrency'

| FILE_NAME        | LAYER_NAME   | TENANT_NAME   | HOST       | SECTION    | KEY             | VALUE |
|------------------|--------------|---------------|------------|------------|-----------------|-------|
| global.ini       | DEFAULT      |               |            | execution  | max_concurrency | 0     |
| indexserver.ini  | DEFAULT      |               |            | execution  | max_concurrency | 10 <--- Effective value
```

> **i 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 value 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

M_INIFILES System View
M_INIFILE_CONTENTS System View
ALTER SYSTEM ALTER CONFIGURATION Statement (System Management)
SAP HANA Configuration Parameter Reference

9.2.3.1.2 Add a System Property Section

Sections in a configuration file bundle parameters of the same category.

Prerequisites

- Your database user has the system privilege INIFILE ADMIN.

- You must have the system privilege DATABASE ADMIN to change the system property of a tenant database from the system database.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Manage Database Configuration* link on the *Database Administration* card.
2. Filter by *Configuration File* in order to display the file you want to add the new section to.
3. Select *Add Section*.
4. Select a file from the dropdown list.
5. Enter the name of the new section.
6. Specify a parameter.
7. Specify the layer to which the parameter applies.
8. Specify:

Layer	Specify
System Layer	Specify the value for the system. The value is also inherited by tenants and hosts unless overridden.
Database Layer	Select one or more databases and specify the value for the parameter.
Host Layer (Can only be used in SYSTEMDB)	Select one or more hosts and specify the value.

9. (Optional) Add a comment about the new parameter.
10. Save the new section.

9.2.3.1.3 Add a System Property Parameter

In the configuration files of an SAP HANA system, you can add a parameter.

Prerequisites

- Your database user has the system privilege INIFILE ADMIN.
- You must have the system privilege DATABASE ADMIN to change the system property of a tenant database from the system database.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Manage Database Configuration* link on the *Database Administration* card.
2. Filter by *Configuration File* and *Section* in order to display the section you want to add the new parameter to.
3. Click the + *(Add parameter to <section_name>)* icon beside the section name.
4. Specify a parameter.
5. Specify the layer to which the parameter applies.
6. Specify:

Layer	Specify
System Layer	Specify the value for the system. The value is also inherited by tenants and hosts unless overridden.
Database Layer	Select one or more databases and specify the value for the parameter.
Host Layer (Can only be used in SYSTEMDB)	Select one or more hosts and specify the value.

7. (Optional) Add comment about the new parameter.
8. Save the new parameter.

9.2.3.1.4 Modify a System Property in SAP HANA Cockpit

In the configuration files of an SAP HANA system, you can modify parameter values.

Prerequisites

- Your database user has the system privilege INIFILE ADMIN.
- You must have the system privilege DATABASE ADMIN to change the system property of a tenant database from the system database.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

You can change the value of a parameter or override its default value. You can't change the name of a section or parameter.

> **i Note**
>
> In general, we don't recommend changing the default values of parameters unless stated in the documentation or instructed by SAP Support. For more information about configuration parameters, refer to the online reference in the SAP Help Portal.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Manage Database Configuration* link on the *Database Administration* card.
2. Filter by *Configuration File* and *Section* in order to display the parameter you want to modify.
   All the parameters in the section are listed. *Override Value* appears for the default layer of each parameter.
   If the default value has been overridden at a level, then the ✎ *(Edit Parameter)* and ⊗ *(Delete Parameter)* icons appear on the layer.
3. To modify a user-defined value on a layer:
   a. Click the ✎ *(Edit Parameter)* icon.
   b. Click in the value field and enter the new value.
   c. Click *Save*.
4. To override a default value, click *Override Value*.
   a. Select one or more layers to apply the new value to.
   b. Enter the required details for the selected layers.
   c. Click *OK* to apply the override.

Related Information

SAP HANA Configuration Parameter Reference

9.2.3.1.5 Restore a System Property Default in SAP HANA Cockpit

You can reset changed parameters to their default values.

Prerequisites

- Your database user has the system privilege INIFILE ADMIN.
- You must have the system privilege DATABASE ADMIN to change the system property of a tenant database from the system database.

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Manage Database Configuration* link on the *Database Administration* card.
2. Filter by *Configuration File* and *Section* in order to display the parameter you want to restore to its default value.
   All the parameters in the section are listed. *Override Value* appears for the default layer of each parameter.
   If the default value has been overridden at a level, the ✎ *(Edit Parameter)* and ⊗ *(Delete Parameter)* icons appear on the layer.
3. Do one of:

To	Action
To restore the default value for a specific layer	Click the ⊗ *(Delete Parameter)* icon beside the layer. The user-defined value is cleared and the default value is reapplied.
To restore the default value for all layers	1. Click *Override Value*. 2. Select the layers to restore. 3. Click *Restore Default for All*. The value for each layer is restored to the default value. 4. Click *OK*.

9.2.3.1.6 View Change History

You can view the change history of one or more values.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

If a row in the table is selected when you click *View Change History*, then the configuration file, section, and parameter are prefilled in the Change History screen and the screen automatically loads the data.

If a row isn't selected when this screen is launched, then only the filters set on the screen that launched it are set and the table is populated.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Manage Database Configuration* link on the *Database Administration* card.
2. While the *Parameters* tab is active, select *View Change History*.
   The change history table is displayed, sorted by time. You can personalize the table, including grouping, sorting, hiding columns, column order.

9.2.3.2 Create a Configuration Template

Set up a configuration template of parameter settings from one database that you can later apply to other databases.

Prerequisites

You're familiar with the layers at which parameters can be configured. See *Database Configuration in SAP HANA Cockpit*.

Procedure

1. On the SAP HANA cockpit *Home* page, click the *Configuration Templates* tile.
2. Select *Create Template*.
3. Enter a name and description for the new template.
4. Specify the layer from which you select parameters for this template.
5. Select a database from which you select parameters for this template.
6. Select parameters for the template.
7. Select *Review*. If necessary, click *Edit* to change your selections.
8. Select *Create Template*

Related Information

Database Configuration in SAP HANA Cockpit [page 443]

9.2.3.3 Apply a Configuration Template

Use a configuration template to copy parameter settings from one database to another.

Prerequisites

- Your database user has the system privilege INIFILE ADMIN.
- A configuration template has been created.

Procedure

1. On the SAP HANA cockpit *Home* page, click the *Configuration Templates* tile.
2. Choose a template.
3. Select *Apply to Databases*.
4. Choose a database to which to apply the template.
5. Click *Apply Template*.

Results

A message tells you whether the template was applied.

Related Information

Create a Configuration Template [page 452]

9.2.3.4 Compare Configurations

You can compare the configuration of on database to that of another or compare the configuration of a database with a snapshot.

Prerequisites

- Your database user has the system privilege INIFILE ADMIN.

- You must have the system privilege DATABASE ADMIN to change the system property of a tenant database from the system database.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Manage Database Configuration* link on the *Database Administration* card.
2. Choose a type and a source.

Compare	Options
System Database	○ You can choose from any system database that you're authorized to view. ○ You can choose one or more databases as targets.
Tenant Database	○ You can choose from any tenant database that you're authorized to view. ○ You can choose one or more tenants as targets.
Snapshots	○ You can choose from snapshots that were created for this database.

3. Click *Compare*.
4. Choose a target.

   The target must be of the same type as the source.
5. (Optional) Click *Show Differences Only*.
6. Enter a configuration file name, section, and layer, or leave them blank to have all the parameters included in the results.

   The list is updated automatically.

9.2.3.5 Snapshots

You can take snapshots, view details, compare, and delete snapshots.

9.2.3.5.1 Take a Snapshot

At any time, you can create a new snapshot.

Prerequisites

- You have the CATALOG READ system privilege.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Manage Database Configuration* link on the *Database Administration* card.
2. Select the *Snapshots* tab, if not active.
3. Click *Take Snapshot*.
4. (Optional) Enter a descriptive name, and select *Take Snapshot*.

Results

The new snapshot is added to the list.

9.2.3.5.2 View Snapshot Details

You can drill down from the list of snapshots to view parameter details of a particular snapshot.

Prerequisites

- You have the CATALOG READ system privilege.

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Manage Database Configuration* link on the *Database Administration* card.
2. Select the *Snapshots* tab, if not active.
3. Click the ❯ icon of the snapshot you want to view.

Results

Details of the selected snapshot appear.

9.2.3.5.3 Delete a Snapshot

You can remove a snapshot from the list of snapshots.

Prerequisites

- You have the CATALOG READ system privilege.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Manage Database Configuration* link on the *Database Administration* card.
2. While the *Snapshots* tab is active, select the snapshot you want to delete.
3. Select the ⊗ *(Delete Snapshot)* icon.
4. Confirm the deletion.

9.2.3.5.4 Compare Snapshots

You can compare different snapshots.

Prerequisites

- You have the CATALOG READ system privilege.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Manage Database Configuration* link on the *Database Administration* card.
2. Select the *Snapshots* tab, if not active.
3. Select the source snapshot you want to compare.
4. Click *Compare*.
5. Select a target snapshot to compare to.
6. Select *Compare*.
7. Click the icon in the *Snapshot* field, and select a target snapshot to compare to.

Results

Only the differences between the two snapshots are listed. To show all values in both snapshots, clear the check mark from *Show Differences Only*.

9.2.4 Managing Workload Classes

Several configuration options are available so that you can tailor workload classes in the SAP HANA database to your needs.

You can manage workload in SAP HANA by creating workload classes and workload class mappings. Workload classes and mappings are SQL object for workload management in SAP HANA. The goal of workload classes and mappings is to provide an easy way for administrators to regulate applications based on pre-defined mapping rules in order to avoid resource shortages with regard to CPU and memory consumption. Appropriate workload parameters are dynamically applied to each client session.

You can classify workloads based on user and application context information and apply configured resource limitations (for example, a statement memory limit). Workload classes allow SAP HANA to influence dynamic

resource consumption on the session or statement level. When a request from an application arrives in SAP HANA, the corresponding workload class is determined based on the information given by the session context such as application name, application user name and database user name. Once the corresponding workload class is determined, the application request can have its resources limited according to the workload class definition.

Statement memory limits will not apply if memory tracking is inactive in SAP HANA cockpit. You can activate memory tracking in the Configuration settings.

Related Information

Create a Workload Class [page 460]
Create a Workload Class Mapping [page 462]
Create User-Specific Parameters [page 465]
Apply Global Settings [page 458]
Disable or Enable a Workload Class [page 466]
Import Workload Classes [page 467]
Export Workload Classes [page 468]

9.2.4.1 Apply Global Settings

You can apply global settings which are used as default values for workload classes. Enabling memory tracking allows you to also monitor the amount of memory used by single workload classes.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Workload Classes lists existing workload classes and provides you with information about the workload handling of the database. You can create and edit workload classes and corresponding workload class mappings.

Procedure

1. On the *Database Overview* page, click the *Manage workload classes* link on the *Database Administration* card.

   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 and thread limits, the total memory and thread limits, and the number of mappings. For databases running SAP HANA SPS 03 or higher, you can also see the query timeout value.

   > **i Note**
   >
   > Not all of the columns listed below are visible by default. You can add and remove columns in the Columns dialog, which you open by clicking the ⚙ *(Settings)* icon in the table toolbar.

2. To monitor the memory consumption of workload classes, enable memory tracking using *Monitor Statements* .

   Information about the memory consumption of workload classes is collected and displayed.

   For more information about memory tracking and setting memory limits, see *Setting a Memory Limit for SQL Statements* in the *SAP HANA Administration Guide*.

3. To limit the memory consumption and number of threads per statement for the system globally select *Edit Global Limits*.
4. Specify values for the following fields:

Field Name	Description
Limit Type	Individual Statement Limit or Total Aggregate Statement Limit
Statement Memory Limit	Displayed if Individual Statement Limit is the specified limit type. Maximum amount of memory the statement may use, in GB.
Total Memory Limit	Displayed if Total Aggregate Statement Limit is the specified limit type. Maximum amount of memory all statements may use, in GB.
Statement Thread Limit	Displayed if Individual Statement Limit is the specified limit type. Maximum number of parallel threads the statement may execute
Total Thread Limit	Displayed if Total Aggregate Statement Limit is the specified limit type. Maximum number of parallel threads all statements may execute.
Query Timeout	The amount of time in seconds before the query times out. (Available for databases running SAP HANA SPS 03 or higher).

5. Select *Save*.
6. 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

Create a Workload Class [page 460]
Create a Workload Class Mapping [page 462]
Create User-Specific Parameters [page 465]
Disable or Enable a Workload Class [page 466]
Import Workload Classes [page 467]
Export Workload Classes [page 468]
Setting a Memory Limit for SQL Statements

9.2.4.2 Create a Workload Class

You can create workload classes to manage the workload of the SAP HANA system.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

You can classify workloads based on user and application context information and apply configured resource limitations (for example, a statement memory limit). Workload classes allow SAP HANA to influence dynamic resource consumption. A workload class must contain at least one workload class mapping that specifies the workload based on user and application context information.

Procedure

1. On the *Database Overview*, click the *Manage workload classes* link on the *Database Administration* card.

   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 and thread limits, the total memory and thread limits, and the number of mappings.

   > **i Note**
   >
   > Not all of the columns listed below are visible by default. You can add and remove columns in the Columns dialog, which you open by clicking the ⚙ *(Settings)* icon in the table toolbar.

2. Specify the workload class details. The fields available are determined by the type and version of your database.

Field Name	Description
Workload Class Name	A name for the new workload class
Execution Priority	Priority, from 0 (lowest) to 9 (highest)
Limit Type	Individual Statement Limit or Total Aggregate Statement Limit
Statement Memory Limit	Displayed if Individual Statement Limit is the specified limit type. Maximum amount of memory the statement can use, as either an absolute or relative value.
Total Memory Limit	Displayed if Total Aggregate Statement Limit is the specified limit type. Maximum amount of memory all statements can use, as either an absolute or relative value.
Statement Thread Limit	Displayed if Individual Statement Limit is the specified limit type. Maximum number of parallel threads the statement can execute, as either an absolute or relative value.
Total Thread Limit	Displayed if Total Aggregate Statement Limit is the specified limit type. Maximum number of parallel threads all statements can execute, as either an absolute or relative value.
Query Timeout	The amount of time in seconds before the query times out. (Available for databases running SAP HANA SPS 03 or higher).
Uncommitted Write Lifetime Limit	The duration of uncommitted write transactions, in minutes, before the connection is terminated.
Idle Cursor Lifetime Limit	The duration of cursors, in minutes, before the connection is terminated.

3. You can also immediately create a mapping for the workload class by entering the mapping properties under *Mapping Details (Optional)*. Refer to *Creating a Workload Class Mapping* for details.
4. Select *Create*.

Results

The workload class is created and displayed in the list. If you have specified mapping properties, a mapping is also created and assigned to the workload class. Before you apply a workload class, you need to assign a mapping.

Related Information

Create a Workload Class Mapping [page 462]
Create User-Specific Parameters [page 465]
Disable or Enable a Workload Class [page 466]
Import Workload Classes [page 467]
Export Workload Classes [page 468]
Create a User Group [page 340]

9.2.4.3 Create a Workload Class Mapping

Mappings link workload classes to client sessions depending on the value of a specific client information property. A workload class must contain at least one workload class mapping that specifies the workload based on user and application context information.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, click the *Manage workload classes* link on the *Database Administration* card.

   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 and thread limits, the total memory and thread limits, and the number of mappings. For databases running SAP HANA SPS 03 or higher, you can also see the query timeout value.

   > **i Note**
   >
   > Not all of the columns listed below are visible by default. You can add and remove columns in the Columns dialog, which you open by clicking the *(Settings)* icon in the table toolbar.

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 and clicking *Create*.
3. To create a workload class mapping, you must enter a *Mapping Name* and specify a value for at least one other field.

   The workload class with the greatest number of matching properties to the session variables passed from the client is applied. If two workload mappings have the same number of matching properties then they are matched in the prioritized order as listed in the table: *application user name* > *client* > *application component name* > *application component type* > *application name* > *user name*. For example, a mapping where the application user matches takes precedence over a mapping where the database user matches (assuming an equal number of matching properties).

   The fields available depend on your database type and version.

Field Name	Description
Mapping Name	Name of the workload mapping.
Schema	Schema name of object defined in the OBJECT NAME property. If you enter a value for *Schema*, then you must enter a value for *Object*.

Field Name	Description
Object	Object types PROCEDURE, PACKAGE and AREA are supported. This property only applies to procedures including AFLLANG procedure which is a standard execution method to execute the application function. Example: If a workload class is matched to an object with type AREA, then it will apply the workload class definition to all AFLLANG procedures which call application functions in the given AFL AREA. Object type PACKAGE works in a similar way. If more than one workload class is matched by the OBJECT NAME then the more specific object type has the higher priority: PROCEDURE > PACKAGE > AREA. If you enter a value for *Object*, then you must enter a value for *Schema*.
Application User Name	Name of the user logged in to the application.
XS Application User Name	Name of the XS application user. For XSA applications which use the session variable XS_APPLICATIONUSER for the business user value.
Client	ABAP client number. For example, 000.
Application Component Name	Name of the application component. This value is used to identify sub-components of an application, such as CRM inside the SAP Business Suite. For example, `/SSB/ALERT_NOTIFICATION_REPORT`.
Application Component Type	Name of the component type. This value is used to provide coarse-grained properties of the workload generated by application components.
Application Name	Name of the application.
Database User Name	Name of the database user. For example, SYSTEM. You cannot enter both a user name and a user group.
User Group Name	Name of the user group. If you wish, you instead of entering a name, you can first create a new user group by selecting the *Add User Group* link. You cannot enter both a user name and a user group.

4. Select *Create*.

Results

The workload class mapping is created and displayed in the list.

Related Information

Managing Workload with Workload Classes
Create User-Specific Parameters [page 465]
Disable or Enable a Workload Class [page 466]
Import Workload Classes [page 467]
Export Workload Classes [page 468]
Create a User Group [page 340]

9.2.4.4 Edit a Workload Class

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, click the *Manage workload classes* link on the *Database Administration* card.

   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 and thread limits, the total memory and thread limits, and the number of mappings.

   > **i** Note
   >
   > Not all of the columns listed below are visible by default. You can add and remove columns in the Columns dialog, which you open by clicking the ⚙ *(Settings)* icon in the table toolbar.

2. Select a workload class entry in the list.
3. Select *Edit*.
4. In the *Edit Workload Class* dialog, modify the desired fields.

Field Name	Description
Workload Class Name	A name for the workload class
Execution Priority	Priority, from 0 (lowest) to 9 (highest)
Limit Type	Individual Statement Limit or Total Aggregate Statement Limit
Statement Memory Limit	Displayed if Individual Statement Limit is the specified limit type. Maximum amount of memory the statement may use, in GB.
Total Memory Limit	Displayed if Total Aggregate Statement Limit is the specified limit type. Maximum amount of memory all statements may use, in GB.
Statement Thread Limit	Displayed if Individual Statement Limit is the specified limit type. Maximum number of parallel threads the statement may execute
Total Thread Limit	Displayed if Total Aggregate Statement Limit is the specified limit type. Maximum number of parallel threads all statements may execute.

5. Select *Save*.

9.2.4.5 Create User-Specific Parameters

User-specific parameters can be created for workload classes.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

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. On the *Database Overview* page, click the *Manage workload classes* link on the *Database Administration* card.

   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 and thread limits, the total memory and thread limits, and the number of mappings. For databases running SAP HANA SPS 03 or higher, you can also see the query timeout value.

   > **i Note**
   >
   > Not all of the columns listed below are visible by default. You can add and remove columns in the Columns dialog, which you open by clicking the ⚙ *(Settings)* icon in the table toolbar.

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
Statement Thread Limit	Maximum number of parallel threads the statement may execute.

Results

The user-specific parameters are created and displayed in the list.

- Note that a user parameter-based approach to limiting memory for statements is not supported for cross-database queries, nor is it effective in XSC developed applications. In these cases you can apply memory limits using workload classes in the remote tenant database.
- Similarly, the user priority value is not effective in XSC developed applications. For XSC applications you can apply a priority value using workload classes.
- 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.

9.2.4.6 Disable or Enable a Workload Class

You can disable or enable workload classes.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

After creating one or more workload classes, you can disable them. This may be necessary for testing purposes. You can also enable workload classes that have been previously disabled.

Procedure

1. On the *Database Overview* page, click the *Manage workload classes* link on the *Database Administration* card.

   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 and thread limits, the total memory and thread limits, and the number of mappings. For databases running SAP HANA SPS 03 or higher, you can also see the query timeout value.

   > **i** Note
   >
   > Not all of the columns listed below are visible by default. You can add and remove columns in the Columns dialog, which you open by clicking the ⚙ *(Settings)* icon in the table toolbar.

2. To disable workload classes:
   a. From the overflow menu above the table, select *Disable*.
   b. Select one or more workload classes.
   c. Select *OK*.

   The workload classes are disabled.
3. To enable workload classes:
   a. Select *Enable*.
   b. Select one or more workload classes.
   c. Select *OK*.

   The workload classes are enabled.

9.2.4.7 Import Workload Classes

Workload classes can be imported from another system, as in the case of going from a test system to a production system.

Prerequisites

You have a file containing workload classes exported from another system.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, click the *Manage workload classes* link on the *Database Administration* card.

   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 and thread limits, the total memory and thread limits, and the number of mappings. For databases running SAP HANA SPS 03 or higher, you can also see the query timeout value.

   > **i** Note
   >
   > Not all of the columns listed below are visible by default. You can add and remove columns in the Columns dialog, which you open by clicking the ⚙ *(Settings)* icon in the table toolbar.
2. From the overflow menu above the table, select *Import*.
3. In the dialog, select *Browse* to specify where to retrieve the file containing previously exported workload class definitions.

4. Use the dialog options to specify what should happen when an imported workload class matches one that is already in the list:
   - Do not import the file
   - Only import classes or mappings that are not duplicates
   - Import the file; overwrite duplicate classes or mappings
   - Remove all classes and mappings before importing the file

   Because Query Timeout is only available on databases running SAP HANA SPS 03 or higher, if the query timeout is a value other than 0, you will not be able to import from a database running SAP HANA SPS 03 or higher to a database running an earlier version of SAP HANA.

Results

The *Import Results* dialog shows the number of workload class definitions that were successfully imported, and may include details about workload class definitions that failed to import or that were overwritten or skipped.

Related Information

Export Workload Classes [page 468]

9.2.4.8 Export Workload Classes

Workload classes can be exported in preparation for importing them into another system, as in the case of going from a test system to a production system.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, click the *Manage workload classes* link on the *Database Administration* card.

   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 and thread limits, the total memory and thread limits, and the number of mappings.For databases running SAP HANA SPS 03 or higher, you can also see the query timeout value.

> **Note**
>
> Not all of the columns listed below are visible by default. You can add and remove columns in the Columns dialog, which you open by clicking the *(Settings)* icon in the table toolbar.

2. Select one or more workload classes.

   Because Query Timeout is only available on databases running SAP HANA SPS 03 or higher, if the query timeout is a value other than 0, you will not be able to import from a database running SAP HANA SPS 03 or higher to a database running an earlier version of SAP HANA.
3. From the overflow menu above the table, select *Export*.
4. Follow the prompts in the dialog to specify where to save the file.

Related Information

Import Workload Classes [page 467]

9.2.4.9 Monitor Workload Classes

You can view monitoring and analysis information on workload class usage.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, click the *Manage workload classes* link on the *Database Administration* card.

   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 and thread limits, the total memory and thread limits, and the number of mappings. For databases running SAP HANA SPS 03 or higher, you can also see the query timeout value.
2. To monitor workload classes and active statements, select *Monitor*.
3. Use the filter to select one or more workload classes.

   The table displays the workload classes. The graph displays the number of active statements, arranged by used memory and number of threads per statement.

4. (Optional) Adjust the data refresh rate using the refresh icon at the top right.
5. (Optional) Select a specific row in order to drill down.

   The display changes to show the active statements for the selected workload class.
6. (Optional) Select *Details* to drill down to statement details.

9.3 Recommendations

Use recommendations to get suggestions for changes you can apply to the SAP HANA database to increase its performance and operation.

In contrast to alerts, recommendations are not warnings that you must react to. They offer rule-based advice on database performance improvement and show the next steps on how to implement the suggested changes.

There are three different types of recommendations:

- *Ad hoc* recommendations are based on an SAP-internal support tool that logs the most impactful changes you can make to improve performance of the SAP HANA database.
- *Physical Design* recommendations target the object design within the database and make suggestions on which changes to it could offer the most positive impact to the SQL performance, memory footprint, CPU usage, etc.
- *SQL Recommendations* are based on analysis of the SQL statements executed by the system and offer suggestions regarding changes to SQL statements.

How can you access recommendations?

You can access recommendations from the Database Overview page of the SAP HANA cockpit as follows:

- Open the *Recommendations* card.
- Open the *SQL Analyzer* tool and go to the *Recommendations* section.

  > **i** Note
  >
  > The *Recommendations* section is only visible when there are any available recommendations.

  To find out how to access the SQL analyzer, see *Analyzing SQL Performance*.

Related Information

View and Follow Recommendations [page 471]
Configure Recommendations [page 472]
Analyzing Statement Performance [page 568]

9.3.1 View and Follow Recommendations

Recommendations is used for suggestions for database performance and operation improvements.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, click the *Administration* view and click the *Recommendations* card.

   The *Recommendations* page opens, displaying the overview page with three different types of recommendations.
   - Ad hoc
   - Physical Design
   - SQL Recommendations.

   Each section displays the recommendation count.
2. Select a recommendation you wish to follow up.

   The detailed recommendation view opens. This information and actions available on this page differ depending on the recommendation type.

Recommendation type	Information	Navigation options
Ad hoc	○ Impact ○ Data source ○ Recommended actions	SAP Notes
Physical Design	○ Impact ○ Data source ○ Recommended actions You can search, sort, and configure the following sections: ○ Related object or object tables, or ○ Related SQL statement strings	○ Suggested application to address the presented issue ○ From a SQL statement string, SQL Analyzer.

Recommendation type	Information	Navigation options
SQL Recommendations	○ Impact ○ Data source ○ Recommended actions You can search, sort, and configure the following section: ○ Related SQL statement strings	○ Suggested application to address the presented issue ○ From a SQL statement string, SQL Analyzer, or ○ Plan Stability

Optional: You can enlarge the detailed recommendation view to fit the screen.

3. Follow up on the recommended action by navigating to the suggested application.
4. Optional: If you do not choose to navigate to the suggested application, you can exit the detailed recommendation view and go back to the overview page or choose another recommendation from the sidebar.

Related Information

Recommendations [page 470]
Configure Recommendations [page 472]
SQL Plan Stability [page 588]
Analyzing Statement Performance [page 568]

9.3.2 Configure Recommendations

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, click the *Administration* view and click the *Recommendations* card.
2. Choose the *Configure* link.

   The *Configure Recommendations* dialog opens. The different types of configurations can be found in four tabs: general, physical design, SQL, and generate recommendations.
3. Choose to show only recommendations collected for the last month, week, day, or two hours in the *General* category.

4. In the *Physical Design* category, you can adjust the triggers for the following recommendations:
    - The threshold for the *Review data statistics* recommendation
    - The threshold for the *Review table distribution status* recommendation
    - The size and percentages of hot and cold object thresholds for the *Update load unit* recommendation.
5. In the *SQL Recommendations* category, you can adjust the triggers for the following recommendations:
    - The maximum and minimum execution count for the *Enable/disable Abstract SQL Plan* recommendation
    - The minimum total result count for the *Add more filters* recommendation
6. Choose which data sources to use for recommendations in the *Generate Recommendations* category.

    You can change the collection status or navigate to the application for further configuration for the following:
    - Data Statistics Advisor
      Optional: You can choose to clear all data statistics recommendations.
    - Resource Tracking for Memory Usage and Network Transfer in Statement Level
    - Plan Stability
    - SQL Analyzer
    - Native Storage Extension (NSE) Advisor
      Optional: You can choose to clear all NSE recommendations.

Related Information

Recommendations [page 470]
View and Follow Recommendations [page 471]
Managing Warm Data With the Native Storage Extension [page 147]

9.3.3 Use the NSE Advisor for Warm Data Recommendations

Use the native storage extension (NSE) Advisor to get suggestions about load units for tables, partitions, or columns according to how frequently they are accessed.

Prerequisites

Depending on the SAP HANA version you are using, you require the following role to perform the operations in this task:

Version	Role
SAP HANA 2.0 SPS 04 **Revisions 40, 41 or 42**	`SAP_INTERNAL_HANA_SUPPORT`

Version	Role
SAP HANA 2.0 SPS 04 **Revisions 43 and higher**	`CATALOG READ` or `DATA ADMIN`

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

The recommendation is generally to change the object to:

- Page-loadable to reduce memory footprint without much performance impact.
- Column-loadable to improve performance by keeping hot objects in memory.

Procedure

1. On the *Database Overview* page, click the *Administration* view and click the *Recommendations* card.

   The SAP HANA cockpit loads the recommendations, with each recommendation having an impact level of High, Medium, or Low.
2. Click on the *Ad Hoc*, *Physical Design*, and *SQL Recommendation* tabs to see detailed analysis for each of these subjects. See *Recommendations* for more information.
3. Use the ALTER TABLE command to make changes to the tables load units. See *ALTER TABLE* in the *SAP HANA SQL Reference Guide*.
4. (Optional) Click the *Configure* button to display the *Configure Recommendations* screen. From this screen, select one of the following:
   - The *General* tab to determine the time period from which to show recommendations.
   - The *Physical Design* tab to determine the thresholds for the following:
     - The ratio between the executed output record and the estimated output size.
     - The average service network request size (in GBs).
     - The minimum size for an object to be considered for a load unit, the percentage of object to consider as hot objects, and the percentage of objects to consider as cold objects.
   - The *SQL* tab to determine the minimum:
     - Execution count and gain ratio for abstract plans.
     - Size of the result set for the threshold of a large result count.
   - The *Generate Recommendations* tab to:
     - Determine if the NSE Advisor collects statistics. On by default. Select *Clear Data Statistics Recommendations* to clear the statistics from the cache.
     - Determine if the NSE Advisor tracks resources for memory usage and network transfers at the statement level. Off by default.
     - Open the *Plan Stability* application. See *Managing Plan Stability* .
     - Open the *SQL Analyzer* in the SAP HANA database explorer.
     - Enable the NSE Advisor to collect statistics. Select *Clear NSE Recommendations* to clear the recommendation cache.

- Analyze Statement Performance. See *Analyze Statement Performance*.

Related Information

Recommendations [page 470]
ALTER TABLE Statement (Data Definition)
SQL Plan Stability [page 588]
Statement Performance [page 570]

9.4 Alerting and Diagnostics

View database information and troubleshoot unresponsive databases.

9.4.1 Collect and Download Diagnosis Information

To help SAP Support analyze and diagnose problems with the SAP HANA database, you can collect diagnosis information into a zip file, which you can then download and attach to a support message for example. With the SAP HANA cockpit, you can create and manage system information dumps.

Prerequisites

- If the database is online, you need the following privileges:

To...	You Need...
Collect diagnosis information	EXECUTE privilege on the procedure `SYS.FULL_SYSTEM_INFO_DUMP_CREATE`
List diagnosis information	SELECT privilege on the view SYS.FULL_SYSTEM_INFO_DUMPS 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.
Download collected diagnosis information	EXECUTE privilege on the procedure `SYS.FULL_SYSTEM_INFO_DUMP_RETRIEVE`

To...	You Need...
Delete collected diagnosis information	EXECUTE privilege on the procedure `SYS.FULL_SYSTEM_INFO_DUMP_DELETE`

- If the system is online, but you want to switch it to offline before collecting information, you will be prompted to connect to the database using the SAP Control credentials.
- If the system is offline (including the system database in a multiple-container system), you must have credentials of the operating system administrator (user `<sid>`adm).
- If the database is a tenant database in a multiple-container system and it is offline, you must be logged on to the system database and have the privileges listed above. It is not possible to collect, list, download, or delete diagnosis information from an offline tenant database.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, select the *Administration* view and select *Manage full system information dumps* on the *Alerting & Diagnostics* card.
2. On the *Diagnosis Files* page, if the system is online, you can use the drop down list to switch to offline. You will be prompted to connect to the database with the SAP Control credentials. (If the system is offline, you cannot switch to online).
3. On the *Diagnosis Files* page, choose a zip file from the list or click *Collect Diagnostics* to create a new zip file.
4. When creating a new zip file, specify the scope of information to be collected:

Option	Description
Collect from existing files	Select this option if you want to collect diagnosis information for one or more file types, for a specific time period, by default the last 7 days. If you also want information from system views, then select *Include system views*. **i Note** If you are connected to the system database of a multiple-container system, only information from the system views of the system database will be collected. Information from the system views of tenant databases will **not** be collected regardless of this option. Information from system views is collected through the execution of SQL statements, which may impact performance. In addition, the database must be online, so this option is not available in diagnosis mode.
Create from runtime environment (RTE)	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: ○ The number sets to be collected (that is, the number of points in time at which RTE dump files will be collected). Possible values are 1- 5. ○ The interval (in minutes) at which RTE dump files are to be collected (possible values are 1, 5, 10, 15, and 30). The default value is 1. ○ The host(s) from which RTE dump files are to be collected.

Option	Description
	○ The service(s) for each selected host from which RTE dump files are to be collected. ○ The section(s) from each selected service from which RTE dump files are to be collected.

5. If your system database is version HANA 2.0 SPS 04 revision 45 or later, then you have the option of choosing which database or databases to collect diagnosis for.

   If the SYSTEM database is online, then you can choose to download diagnostic information for **only** the SYSTEM database, for specified tenant databases, or for all databases in the system.

   If the SYSTEM database is offline, then you can choose to download diagnostic information for either the SYSTEM database or for specified tenant databases.

Results

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.

9.5 Database Information

Accessing general information about the SAP HANA database, such as operational status and database version, can assist you to monitor your database.

In SAP HANA cockpit, you can access *Database Information* by drilling down in the *Database Overview* when the *Administration* or *All* view is selected. To do this, your database user needs the system privilege CATALOG READ.

Details include:

- Information such as operational status, database usage type, whether the database has multiple hosts, the number of hosts (if distributed), and database version
- The SAP HANA version history
- Information about the plug-ins that are installed
- The status of replication from your production database to a secondary database. This information is only available and applicable if you are operating a secondary instance of your database (for example, in a high availability scenario). If this is the case, then content from the primary or production instance of your database is replicated to the secondary instance.

9.6 Manage Hadoop Clusters

You can navigate from the SAP HANA cockpit to the Apache Ambari Web site and monitor Hadoop clusters.

Context

You can find more information about Hadoop clusters in the *SAP HANA Administration Guide*.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, clickthe *Manage Hadoop cluster* on the *Other Administration* card.
2. Click + *Add Hadoop Cluster* and enter the cluster name and the Ambari URL (for example, `http://my.ambari.server.url:8080`).

   You are directed to the Ambari Web site.

   The cluster name and URL information that you enter is not saved in the SAP HANA database. When you close the browser, the information is deleted.

Related Information

Using the Database Overview Page to Manage a Database [page 109]

9.7 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.

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.

Related Information

Monitor Remote Statements and Connections [page 479]
SAP HANA Smart Data Access

9.7.1 Monitor Remote Statements and Connections

View detailed information about the remote statements executed and remote connections active in the database.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Use the monitoring tools to monitor:

Remote connections active in the database	Provides details about the connections that were opened in the current session, including when the connection was opened, how many remote statements were executed, and the name of the remote source.
Remote statements executed in the database	Allows you to see the full SQL text of the SQL statements executed on remote sources. It also shows you when the query was started, how long the query took, and the number of records that were returned.

Procedure

1. On the *Database Overview*, scroll to *SAP HANA Smart Data Access Administration*.
2. Choose the information type to monitor.
   - Running Statements
   - Active Connections

Results

Information available for remote connections

Detail	Description
Connection	Specifies the connection ID
Adapter	Specifies the name of the adapter for the remote source
Status	Specifies the status of the connection. Valid entries are CONNECTED or DISCONNECTED
Source Name	Specifies the remote source name
Source User	Specifies the user name on the remote source
Start Time	Start time of first query execution
Statements	Specifies the statement start time
Details	Specifies information on the adapter properties.

Information available for remote statements

Detail	Description
SQL Statement	Specifies the statement string
Start Time	Specifies the statement start time
End Time	Specifies the statement end time
SAP HANA cockpit: Statement Runtime (Seconds)	Specifies the query execution time
Status	Query execution status: • Analyzing: Query is being analyzed by the query optimizer • Optimizing: Query is being optimized by the query optimizer • Executing: Query is running • Closed: Query has completed • Failed: Query execution failed
Rows	Specifies the number of rows returned in the query result
Fetched Size	Specifies the byte size of fetched records
Remote Source Name	Specifies the remote source name
User	Specifies the remote source name
Transaction	Specifies the transaction ID

9.8 Table Distribution

Use the SAP HANA cockpit to manage table distribution. You can view and save the current table distribution, automatically generate an optimized table redistribution, re-run a previously executed plan, or restore a saved plan.

Context

In a scale-out system, tables and table partitions are distributed across multiple hosts. The location of the tables and partitions can affect performance when queries need to access several tables. You may want to redistribute the tables or partitions to different hosts to better optimize particular capabilities. Or you may want to add a new host to the scale-out system and therefore need to redistribute the tables so that some will reside on the new host.

You must have the RESOURCE ADMIN system privilege to see the *Table Distribution* card. Additional privileges may be needed to perform some table distribution tasks.

Related Information

Redistributing Tables in a Scaleout SAP HANA System
Scaling SAP HANA

9.8.1 Table Placement Rules

Specify rules for where new tables/partitions are placed or where tables/partitions are moved during a redistribution.

Because tables interact with and depend on other tables, in a distributed database landscape time can be lost when tables or dependent parts are placed on different hosts. Also, the way table partitions are placed in your landscape could cause critical performance issues. You can classify each table with group type and subtype information and store a set of configuration values for different table classification patterns. During operations like table redistribution, these patterns are used to determine where tables or table partitions are placed (for example, the number of table partitions allowed on a specific host). See also the *Table Placement* section of the *SAP HANA Administration Guide*.

For example, if you create the following table placement rule:

Location	Slave
Repartition when number of rows exceeds	40,000,000
Maximum number of rows per partition	30,000,000

Then:

- If the landscape has 6 slave nodes, and a table has 1 exiting partition and 39,000,000 records, when the placement rule is applied the table will stay at 1 partition, because the 40,000,000 value was not exceeded.
- If the landscape has 6 slave nodes, and a table has 1 exiting partition and 40,000,001 records, when the placement rule is applied the table will get 3 partitions, because the 40,000,000 value was exceeded, and the number of records in the table divided by 30,000,000 was greater than 1.
- If the landscape has 6 slave nodes, and a table has 1 exiting partition and 90,000,001 records, when the placement rule is applied the table will get 6 partitions, because the 40,000,000 value was exceeded, and the number of records in the table divided by 30,000,000 was greater than 3.

9.8.1.1 Add or Edit a Table Placement Rule

Create a new or edit an existing table placement rule.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- You have the TABLE ADMIN system privilege.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the ••• *(More)* icon on the *Table Distribution* card, and then select *Edit Table Placement Rules*.
2. Choose:

Option	Action
Add a new table placement rule.	Select *Add*.
Edit an existing table placement rule.	Click the ✎ *(Edit)* icon for the rule.

3. Specify the tables the rule will apply to.
4. Specify the rule conditions.

Rule Condition	Description
Location	Available locations are: ○ *master*: represents the master node

Rule Condition	Description
	○ *slave*: represents all slave nodes that belong to the worker group 'default' ○ *all*: represents all nodes that belong to the worker group 'default', (master and slave nodes). Click the ℹ *(Information)* icon to see the hosts associated with the selected location. You can define a custom location by clicking *Manage locations*. See *Manage Table Placement Rule Locations*.
Set Persistent Memory	The tables specified in the table placement rule are placed in persistent memory.
Set Page Loadable	A subset or page of the table or partition can be loaded as needed.
Set Replica Count	The minimum number of replicas for replicated tables. This is useful for parallelizing high loads.
Repartition when number of rows exceeds	The minimum number of records that must exist in the table before a calculation of repartitioning begins. (MIN_ROWS_FOR_PARTITIONING)
Maximum number of rows per partition	If the row count in one of the partitions exceeds this value then further splits are considered. (REPARTITIONING_THRESHOLD). The maximum number of partitions for a table is 12, by default. (If necessary, you can change the default by modifying the configuration parameter `max_partitions`. The maximum number of partitions for a table is also limited to the number of available hosts for the specific table as provided by the table placement landscape. For example, if a table is configured by table placement to be located on slave indexserver and the landscape has 3 slave nodes, you will get 3 or fewer partitions for this table. (If necessary, you can change this behavior by modifying the configuration parameter `max_partitions_limited_by_locations`.
Number of hosts to contain data	How many hosts should contain data.
When partitioning, ensure an initial number of partitions	The total number of partitions is the integer double of this value. For example, entering a value of 3 leads to a partition number sequence of 1, 3, 6, 12. (INITIAL_PARTITIONS)
Each table in this group has the same number of partitions	Checking the box specifies that all partitions of the tables in a group will contain the same number of partitions
Split the dynamic others partitions when the number of rows exceeds ...	Applies to tables that use the dynamic range partitioning feature.

5. Select *Save*.

Related Information

Table Placement

9.8.1.2 Copy a Table Placement Rule

Create a new table placement rule by copying an existing rule and reusing the same or modified parameters.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- You have the TABLE ADMIN system privilege.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the ••• *(More)* icon on the *Table Distribution* card, and then select *Edit Table Placement Rules*.
2. Click the *(Copy)* icon for the rule to copy.
3. (Optional) Modify the rule conditions as desired.
4. Select *Save*.

9.8.1.3 Manage Table Placement Rule Locations

Manage a list of locations for table placement rules.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- You have the TABLE ADMIN system privilege.

Context

Custom locations are defined by using the volume IDs of hosts included or excluded in the location.

You can edit the predefined locations *all* and *slave*, but not *default* or *master*.

You may wish to edit the *all* location if it has been used in many table placement rules and you then add extension nodes to the system. The extension nodes will be added to *all*; however, extension nodes should not be used for hot data. You can either change all the rules that refer to *all* to refer to a newly defined custom location (for example, *all2*), or you edit *all* to exclude the extension nodes.

You can also add new locations. The *Manage Locations* option is only available when editing or adding a table placement rule, but you don't have to save the rule to save the new or modified location settings.

You can change the hosts included and excluded from a location but you cannot change the name of an existing custom location.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the ••• *(More)* icon on the *Table Distribution* card, and then select *Edit Table Placement Rules*.
2. Click the ✎ *(Edit)* icon for an existing rule or click the *Add* button
3. In the *Rule Conditions* column, select *Manage Locations*.
4. Choose:

Option	Action
To edit an existing location.	1. Click the location you want to edit. 2. Specify which hosts you want the location to include or exclude. If the selected location is not editable (for example, Default), a list of the hosts for the location appears instead of the include and exclude fields.
Add a new location.	1. Click *Add Location*. 2. Give your location a name. 3. Click the ⧉ icon on the *Include* or *Exclude* fields to select volumes and pre-defined locations to include and exclude from the new location.
Delete a location.	Click the ⊗ *(Delete)* icon beside the location. If the location is used in existing table placement rules, a dialog box appears listing those rules. You must select another location before the *Update Rules and Delete Location* button becomes available and the location is deleted. All of the rules listed are assigned the new location you specify.

9.8.1.4 View Tables for a Select Rule

View those tables impacted by a specific rule.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.

Context

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the ••• *(More)* icon on the *Table Distribution* card, and then select *Edit Table Placement Rules*.
2. Select a rule.
3. Click *View Tables for Selected Rule*.

   The *Current table Distribution* page opens, automatically filtered to only display tables defined in the rule.

9.8.2 Designing Partitions

There are a number of factors to consider to optimize the design of your data partitioning strategy including how it will affect select and insert performance and how it will adjust to data changes over time.

Different partitioning strategies need to be tested to determine the best one for your particular scenario. Based on your tests you should choose the partitioning strategy that shows the best performance for your scenario. The design principals listed here are aids to help you decide on the correct partitioning strategy for your scenario.

> **i** Note
>
> SAP Business Warehouse on SAP HANA handles partitioning itself. Do not interfere with its partition management unless this has been recommended by SAP.

Query Performance

- For replicated dimension tables the database tries to use replicas that are local to the fact table partitions.
- Partition pruning analyzes the WHERE clauses and seeks to reduce the number of partitions. Try to use partitioning columns that are often used in WHERE clauses. This reduces run time and load.
- Usually hash partitioning is the best partitioning scheme for the first level, especially in scale out scenarios. This is because the client may already use pruning on the client machine and send the query directly to the host that holds the data, where possible. This is called "client-side statement routing". This is especially important for single select statements.
- Use as many columns in the hash partitioning as required for good load balancing, but try to use only those columns that are typically used in requests. In the worst case only single select statements may leverage pruning.
- If tables are joined with each other, it is beneficial if the tables are partitioned over the same columns and have the same number of partitions. This way the join may be executed locally in scale out scenarios and the network overhead is reduced.
    - This guarantees that the matching values are in a partition with the same part ID. You have to put all parts with the same ID on the same host.
- Queries do not necessarily become faster when smaller partitions are searched. Often queries make use of indexes and the table or partition size is not significant. If the search criterion is not selective though, partition size does matter.

Data Manipulation Language (DML) Performance

- If insert performance is key to your scenario, a larger number of partitions might show better results. On the other hand, a higher number of partitions may reduce query performance.
- Partition pruning is used during DML operations.
- For replicated column store tables, all DML operations are routed through the host with the master partition (where the replica with Part ID 1 is located).
- If there is a unique constraint on a non-key column, the performance will suffer exponentially with the number of partitions on other servers. This is because the uniqueness on all partitions has to be checked. Therefore, if partitioning is required, consider a low number of partitions and ideally put all partitions on the same host. This way the number of remote calls is reduced.

Data Lifecycle

If time-based partitioning is suitable for the dataset being partitioned, it should always be used as it has a number of advantages:

- The runtime of a delta merge is dependent on the size of the main index. This concept leverages the fact that new data is inserted into new partitions whereas data in old partitions is infrequently updated. Over time, the formerly new partitions become old and new partitions are being created. Delta merges on old partitions are not required anymore. This way the overall runtime of delta merges does not increase with the table size, but remains at a constant level. Using time-based partitioning often involves the use of hash-range partitioning with range on a date column. This requires knowledge of the actual values for range partitioning.

- By using explicit partition management, new partitions can be created, for example, one partition per calendar week and old partitions may be dropped entirely rather than deleting individual rows.
- If you split an index, always use a multiple of the source parts (for example 2 to 4 partitions). This way the split will be executed in parallel mode and also does not require parts to be moved to a single server first.
- Do not split/merge a table unless necessary. These operations write all data into the log which consumes a high amount of disk space. Moreover, the operations take a long time and locks the table exclusively (only selects are allowed during partitioning operations). ADD PARTITION can be used to add additional partitions. If there is no 'others' partition, this call only creates a new partition which is fast and happens in real time after an exclusive lock of the table was acquired. On the other hand, if the table has an others partition, a call to ADD PARTITION causes the existing others partition to be split into a new others partition and newly requested range. This is a costly operation. Therefore it is recommended, that if ADD PARTITION is used frequently in a scenario, the table shall not have an others partition.

Partition Size

- Due to the high number of factors to be considered when evaluating a partitioning scheme a recommendation for partition sizes cannot be provided. If you do not know if you will partition a table at all or with how many partitions you need to start with measurements. Here are some suggested starting points:
  - If a table has less than 500 million rows, do not partition it at all unless:
    - The corresponding tables in joins are partitioned. If they are try to find mutual partitioning columns.
    - Table growth is expected. Since re-partitioning is time consuming, it is recommended to split a table while it is still small.
  - If your table has more than 500 million rows, choose 300 million per partition.
  - Keep in mind that a partition must not exceed 2 billion rows.
- Be aware that a higher number of partitions might lead to higher memory consumption as each partition has its own exclusive dictionary, which is not shared. If each partition stores disjunctive values, this is not an issue. On the other hand, if each partition has similar or the same values this means that the dictionaries have similar data which is stored redundantly. In this case consider using fewer partitions

Table Design

- If the data is replicated into the SAP HANA database, it might be fair from a data consistency perspective to remove a primary key or to extend the key (since the key constraint is enforced in the source database). This way you might be able to have multiple tables with the same partitioning columns even though the original database design would not have allowed it. Having the same partitioning columns is ideal as related data may reside on the same physical host and therefore join operations may be executed locally with no or hardly any communication costs.
- When designing database schemas for dependent hosts, for example, a database structure for business objects with header and leaf nodes, do not use a single GUID column as the primary key. In such a case it is hardly possible to have all related data (for example, a business object instance) on the same host. One option might be to have a GUID as the primary key in the header table and each host, irrespective of its level, could have that GUID as the first primary key column.

- Do not define a unique constraint on a partitioned table unless absolutely necessary.
- On the second partitioning level, a non-primary key column may be used. Still, the unique constraint has to be enforced on all parts of the respective first-level partition. Since all parts of one first-level partition are moved as a whole, this unique check is always local.
- In case the database table is replicated from another database, an others partition for range is generally recommended. If a proper range is not defined, the insert statement will fail and the data will not get replicated properly.
- Ideally tables have a time criterion in the primary key. This can then be used for time-based partitioning. Number ranges and so on can also be used. The advantage of number ranges is that it is easy to form equally sized partitions, but on the other hand it introduces an administrative burden the amount of data that is loaded needs to be closely monitored and new partitions need to be created in advance. In case of actual dates, you only need to periodically create new partitions, for example, before a new quarter starts.

Other Considerations

- Use GET_NUM_SERVERS() in scripts for hash and round-robin partition specifications. This way Table Placement is used to calculate the number of partitions that will be used in the given landscape.
- If it is likely that a table has to be re-split in future and range partitioning is used, define an others partition. (If it is not defined upon table creation, it can be created afterward and if required dropped after the split operation).
- To check whether a table is partitioned, do not consider the existence of a partition specification in the metadata. Instead check IS_PARTITIONED in M_TABLES or for the existence of parts, for example in M_CS_TABLES. It is allowed that a partition specification is defined which does not immediately lead to a partitioned table.

9.8.3 View or Modify a Table Distribution

To support the analysis and monitoring of performance issues in a distributed SAP HANA system, you can see how tables are distributed across the hosts.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.

Context

In the case of partitioned tables, you can see how the individual partitions and sub-partitions are distributed, as well as detailed information about the physical distribution, for example, partition ID, partition size.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and how they are grouped.
3. (Optional) Select the *Analysis* tab to display the number of records, partitions and table sizes per host.
4. Select *Go*.
   A list of tables is displayed.
5. Hover over a value in any column. If the value turns blue, left-click the value to display information about the table or value and additional options that can be performed to manage the value. In the host columns, hover within the table row and then left-click the blue icon to display the information and additional options.

9.8.3.1 View Memory Usage

View available memory usage data for a table or a group of tables.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- You have the SELECT privilege on the schema _SYS_STATISCTICS.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the table to view.
4. Choose:

Option	Task
To see memory usage for all tables in a group,	Left-click a group name.
To see memory usage for a singe table,	Left-click a table name.

5. Select *View Memory Usage*.

6. Click *Yes* to confirm.

9.8.3.2 Load or Unload Tables

Load or unload tables or a group of tables into memory.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- One of:
  - You created the table.
  - You have the UPDATE privilege on the table or table schema.

Context

As the SAP HANA database automatically manages the loading and unloading of tables it is not normally necessary to manually load and unload individual tables and table columns. However, this may be necessary to:

- Precisely measure the total or "worst case" amount of memory used by a particular table (load).
- Actively free up memory (unload).

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, expand the group containing the table to be loaded or unloaded.
4. Choose:

Option	Task
To load all tables in the group,	Left-click a group name and select *Load*. Tables in the group that are already loaded are ignored.
To unload all tables in the group,	Left-click a group name and select *Unload*. Tables in the group that are not loaded are ignored.
To load a singe table,	Left-click a table. Depending on the status of the table, select *Load* or *Unload*.

5. Click *Yes* to confirm.

Related Information

Load/Unload a Column Table into/from Memory
Memory Sizing [page 408]
Memory Management in the Column Store
LOAD Statement (Data Manipulation)
UNLOAD Statement (Data Manipulation)
View or Modify a Table Distribution [page 489]

9.8.3.3 Move a Table to Another Host

Move a table to another host.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- One of:
  - You created the table.
  - You have the ALTER privilege on the table or table schema.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the table to move.
4. In the column of the target host for the move, hover on the table row, and left-click the blue icon.
5. Choose:

If	Do
The table resides on the selected column host,	1. Select *Move Table Here*.

If	Do
The table resides on a host other than the selected column host,	1. Select *Move Table to Another Host* appears. 2. Click the ⌄ icon an select the target host. 3. Click *Move*.

6. Click *Yes* to confirm the move.

9.8.3.4 Move Table Partitions to Another Host

Move table partitions to another host.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- One of:
  - You created the table.
  - You have the ALTER privilege on the table or table schema.
- You are familiar with partitioning concepts and operations, as outlined in the *Table Partitioning* section of the *SAP HANA Administration* guide.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the table partitions to move.
4. In the column of the target host for the move, hover on the table row, and left-click the blue ⇥ icon.
5. Choose:

If	To	Do
The table resides on the selected column host, select *Move Table Partitions Here*.	Move all partitions to the target host,	1. Select the host name.
	Move select partitions to the target host,	1. Expand the host name. 2. Select the partitions to move.

If	To	Do
The table resides on a host other than the selected column host, select *Move Table to Another Host*.	Move all partitions to the target host.	1. Select the target host name from the list. 2. Click *Select All*.
	Move select partitions to the target host.	1. Select the target host name from the list. 2. Select the partitions to move.

6. Click *Move* and then *Yes* to confirm the move.

9.8.3.5 Partition a Table

Partition a table at a single or multilevel.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- One of:
  - You created the table.
  - You have the ALTER privilege on the table or table schema.
- You are familiar with partitioning concepts and operations, as outlined in the *Table Partitioning* section of the *SAP HANA Administration* guide.

Context

You can partition a table using a single or multilevel partitioning schema based on hash, range, or round robin. For information on partition types and their unique benefits, see the *Table Partitioning* section in the *SAP HANA Administration Guide*.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the table to partition.

4. Left-click a table name to partition, and then click *Partition Table...*.
5. Choose the first level partition type.
6. Enter the partition specifications and click *Next*.
7. (Optional) Select *Second-Level Partitioning* and enter the partition specifications.
8. Click *Partition* to complete the table partitioning process.

9.8.3.6 Unpartition a Table

Merge all partitions into one table.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- One of:
  - You created the table.
  - You have the ALTER privilege on the table or table schema.

Context

The data from each partition is merged into a single table.

If you do not want to merge all partitions into one table, then see *Repartition a Table*. To see the current partitioning schema of a table, see *View the Partitions of a Table*

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the table to unpartition.
4. Left-click a table name to unpartition, and then click *Partition Table...*.
5. Click *Merge Partitions*.

Related Information

View the Partitions of a Table [page 496]
Repartition a Table [page 496]

9.8.3.7 View the Partitions of a Table

Display the current partitioning schema definition of a table.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the partitioned table to view.
4. Left-click a partitioned table name.
5. Scroll to the *Partition Definition* section a the bottom of the dialog box.

9.8.3.8 Repartition a Table

Adjust the single or multilevel type and structure of a partitioned table.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- One of:

- You created the table.
- You have the ALTER privilege on the table or table schema.

Context

You change the type and definition of a partition a table by redefining the partitioning schema. If you want to merge all partitions into one table, see *Unpartition a Table*. To see the current partitioning schema of a table, see *View the Partitions of a Table*. For information on partition types and their unique benefits, see the *Table Partitioning* section in the *SAP HANA Administration Guide*.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the table to repartition.
4. Left-click a partitioned table name and then click *Partition Table...*.
5. Choose the first level partition type.
6. Enter the partition specifications and click *Next*.
7. (Optional) Select *Second-Level Partitioning* and enter the partition specifications.
8. Click *Partition* to complete the table partitioning process.

Related Information

Unpartition a Table [page 495]
View the Partitions of a Table [page 496]

9.8.3.9 Perform a Delta Merge

You can merge the column store table's or group of tables' delta storage into a table's main storage.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

- You have the RESOURCE ADMIN system privilege.
- One of:
  - You created the table.
  - You have the UPDATE privilege on the table or table schema.

Context

It may be necessary or useful to trigger a merge operation 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 stores 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.

Unlike a table partition merge, a delta merge does not merge all partitions into one table.

See also *The Delta Merge Operation* in *SAP HANA Administration Guide*.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the table to delta merge.
4. Left-click a table or group name.
5. If a delta merge is possible, *Delta Merge* appears on the menu. Select it.
6. Select *Yes* to confirm.

Related Information

The Delta Merge Operation

9.8.3.10 Truncate a Table

You can truncate a table or group of tables.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- One of:
  - You created the table.
  - You have the DELETE privilege on the table or table schema.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the table to view.
4. Choose:

Option	Task
To truncate tables in a group,	Left-click a group name.
To truncate a single table.	Left-click a table name.

5. Select *Advanced Operations*, then select *Truncate Table*.
6. If you are truncating by group name, select the tables in the group to truncate.
7. Click *Yes* to confirm.

9.8.3.11 Drop a Table

You can drop a table or group of tables.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

- One of:
  - You created the table.
  - You have the DROP privilege on the table or table schema.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the table to view.
4. Choose:

Option	Task
To drop tables in a group,	Left-click a group name.
To drop a single table.	Left-click a table name.

5. Select *Advanced Operations*, then select *Drop Table*.
6. If you are dropping by group name, select the tables in the group to drop.
7. Click *Yes* to confirm.

9.8.3.12 Display Table Content, Metadata, Access Statistics, or Runtime Data

You can display the first 100 rows of data, the XML-formatted meta data, the table access statistics, and the runtime data for any table.

Context

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- To view table content, one of:
  - You created the table.
  - You have the SELECT privilege on the table or table schema.

If you are viewing a partitioned table from SAP HANA version SPS 03 or later, you can also view the partition statistics.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.
2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the table to show.
4. Left-click a table name.
5. Choose:
   - *Show Content*
   - *Show Meta Data*
   - *Show Runtime Data*
   - *Show Access Statistics*

   A new window displays the details.
6. Use the different tabs to move between the displayed information.
7. (Optional) On the *Meta Data* tab, select *Save As* to save the meta data as a local file.
8. (Optional) On the *Runtime Data* tab, to display size information, close the window, select *Load Table* from the popover, and then re-select *Show Runtime Data*.
9. (Optional) On the *Access Statistics* tab, select *Reset* to restart the collection of access statistics data.

9.8.3.13 Export as CSV

You can export a table or a group of tables as a CSV.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- You have the EXPORT system privilege.
- One of:
  - You created the table.
  - You have the SELECT privilege on the table or table schema.
- To export to a client, you also need the CREATE TEMPORARY TABLE system privilege.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *View Current Table Distribution* on the *Table Distribution* card.

2. (Optional) Use the filtering options and display options to refine the list of tables displayed and then click *Go*.
3. If *Show grouped* is enabled, in the *Name* column, expand the group containing the table to export.
4. Choose:

Option	Task
To export all tables by group,	Left-click a group name.
To export a single table.	Left-click a table name.

5. Select *Export as CSV*.
6. Specify the path where the CSV file will be saved.
7. Specify export options by selecting or de-selecting checkboxes.

   Refer to the *EXPORT Statement (Data Import Export)* topic in the *SAP HANA SQL and System Views Reference* for more information on these options.
8. Select *Save*.

9.8.4 View Redistribution Execution History

At any time, you can view previously-executed table redistribution plans.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

The executed table distribution operations displayed include those that are:

- Running
- Finished
- Failed
- Canceled

> **i** Note
>
> Table distribution errors that are Finished include those that are Finished with Errors. Take note of the number of failed steps in the *Finished (Failed)* column.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the ... *(More)* icon on the *Table Distribution* card, and select *View Redistribution Execution History*.
2. Click the > icon beside a table distribution to see the table redistribution details.
3. Choose the *Plan Steps*, *Failed Operations*, or *Parameters* tabs to view specific information.
4. Click the > icon on each row to see further details.

9.8.4.1 Save the Current Table Distribution

You can save a current table distribution as a distribution plan provided that no table distribution operations are currently running.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.

Context

Changing how tables are distributed across the hosts of a distributed SAP HANA system is a critical operation. Therefore, before executing a redistribution operation, it is strongly recommended that you backup the landscape so that it can be restored if necessary.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the ... *(More)* icon on the *Table Distribution* card, and select *View Redistribution Execution History*.
2. Select the *Save Current Table Distribution* button.

   > → Tip
   >
   > This button is not available for selection if an executed table distribution is currently running.

3. Select *Save*.

   The current table distribution is saved as a distribution plan.

9.8.4.2 Restore a Saved Table Distribution Plan

You can restore a distribution plan.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.

Context

Changing how tables are distributed across the hosts of an SAP HANA system is a critical operation. You may need to restore the table distribution from a previous point in time. The *Restore Saved Distribution* button is only available when you select a *Save Current table Distribution* operation.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the ••• *(More)* icon on the *Table Distribution* card, and select *View Redistribution Execution History*.
2. Select a saved current table distribution to restore.
3. Click the *Restore Saved Distribution* button.
4. Click *Restore Saved Distribution* to confirm the restore.

9.8.4.3 Rerun a Table Distribution Plan

You can run a previously-executed table distribution plan .

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.

Context

The *Rerun Plan* button is unavailable if you select a *Save Current Table Distribution* operation.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the ••• *(More)* icon on the *Table Distribution* card, and select *View Redistribution Execution History*.
2. Select a specific table distribution to rerun.
3. Select the *Rerun Plan* button.
4. Specify whether you want to rerun the entire plan or only the failed steps.
5. (Optional) While the plan is running, you can change the number of operations that can be executed simultaneously. Select *Edit0* beside *Maximum Number of Parallel Plan Items per Active Indexserver*.

9.8.4.4 Automatic Table Distribution Analyzer

Configure the behavior of the *Table Distribution Analyzer*.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- You have the INIFILE ADMIN system privilege.

Context

When enabled, the decision as to whether an optimization is necessary is made by an expected minimum improvement on one of the active optimization targets (configurable in the system configuration). Improvement means a reduction by the threshold percent of the spreading on a target KPI for a workergroup (for example, all, slaves, etc.). Modifying the parameters will not interfere with any currently executing operations since operations are executed following a previously generated plan.

A generated redistribution plan must exist before the analyzer can be enabled. The current status of the analyzer appears in the heading on the *Table Redistribution Execution History* page.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the ••• *(More)* icon on the *Table Distribution* card, and select *View Redistribution Execution History*.
2. To enable the *Automatic table distribution analyzer*, click *Edit*.
3. Click the checkbox and enter a frequency and percentage for the analyzer to use.
4. (Optional) Click *Modify Parameters* to further configure the analyzer.

9.8.5 Generate and Execute a Table Redistribution Plan

Redistribution is a two stage process: you generate a redistribution plan, and then execute it.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.

Context

A table redistribution plan is temporary. Once the session is closed, the plan is removed. The most time consuming part is gathering all of the information about the existing landscape. You may wish to generate a plan, analyze the planned outcome, and regenerate a new plan using information that was previously gathered. See also *Redistributing Tables in a Scaleout SAP HANA System* in the *SAP HANA Administration Guide*.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the ••• *(More)* icon on the *Table Distribution* card, and select *Table Redistribution Plan Generator*.
2. Identify the goal of the table redistribution plan.

Goal	Description
Balance table distribution	The load on a scale-out system changes over time with the usage of the system. This option generates a plan to move tables and partitions to their proper hosts if they are currently on invalid hosts according to the rules specified in the TABLE_PLACEMENT table. The plan will check whether a split or merge is necessary

Goal	Description
	and calculates optimal positions for the parts and tables. All types of tables and parts can be moved. However, only the tables that you have permission to view as catalog objects will be affected.
Redistribute tables after adding host(s)	After adding one or more worker hosts to a scale-out system, you may need to redistribute the tables across the active indexservers. This option checks whether new partitions can be created and generates a plan to move the tables and table partitions as necessary.
Housekeeping	Some regular operations need to be done from time to time. This option allows you to perform various operations in the system, such as, optimize compressions, defrag, load table, merge delta. Only the tables that you have permission to view as catalog objects will be affected. Also, you must have appropriate privileges to perform specific housekeeping operations, such as delta merge.
Check the number of partitions	This option evaluates whether or not partitioned tables need to be repartitioned. The plan will specify how partitioned tables will be repartitioned (split or merge) and how newly-created partitions will be distributed. Note that this is only relevant for column-store tables. System tables, temporary tables, and row-store tables are not considered. In a scale-out system, partitioned tables are distributed across different index servers. The location of the different partitions can be specified manually or determined by the database when the table is initially partitioned. Over time, this initial partitioning may no longer be optimal, for example, if a partition has grown significantly.
Check the correct location of tables and partitions	This option generates a plan to move tables and partitions to their proper hosts if they are on invalid hosts according to the rules specified in the TABLE_PLACEMENT table. Only the tables that you have permission to view as catalog objects will be affected.

3. Under *Options*, specify which tables to consider in the redistribution by setting any combination of options:
   - schema(s)
   - table group(s)
   - group type(s)
   - group subtype(s)
   - table name
   - tables with or without LOB files, or both
   - loaded or unloaded tables, or both
   - filled or empty tables, or both
   - used or unused tables, or both

   You can also specify which outcome you are interested in (for example, balance the workload per host), and how much weight you want to apply to that desired outcome. You can specify the order in which to process the tables, and other options.
4. (Optional) If you have selected *Balance table distribution* or *Redistribute tables after adding host(s)* as a goal, you can opt to analyze table groups by turning on the corresponding switch.
5. If you enabled *Table Group Analysis*,
   a. (Optional) Allow grouping by row store tables or temporary tables.

b. (Optional) Enable *Splitting Groups* and define how to split the groups, by table replication or by execution count. Indicate which groups to include and exclude in the split.
c. (Optional) Select *Grouping Options for S4 System* to enable additional S4 specific options when adding or editing an *Analyze Existing Groups* analysis type.
d. Modify an existing analysis or click *Add Analysis* to create a new one.
e. Specify the analysis type to use for the plan.
f. Enter a *Preview* name.
g. Select an analysis source.
h. (Optional) Create an analysis filter.
i. Select a table source.
j. (Optional) Add tables and schemas to the analysis.

6. Click the *Preview Table Groups* button to see the proposed group redistribution.
7. (Optional) Specify any advanced options.
8. Click the *Review* button.
9. Click the *Generate Table Redistribution Plan* button.

   The progress status displays on screen. When progress is complete, the *Table Redistribution Plan* page displays.
10. (Optional) In the *Plan Steps* tab, select a row to review details of a single operation or step group.
11. (Optional) In the *Analysis* tab, use the drop down variant selector, and the filter check boxes in order to display relevant planned and actual data.
12. (Optional) Select *Regenerate with Modified Parameters*.
    a. Indicate whether the table redistribution plan should use landscape information collected from the previously generated plan, or perform a new collection of landscape information.
    b. Adjust the goal of the table redistribution plan, as desired.
    c. Adjust the plan options, as desired.
    d. Select the *Review* button.
    e. Select the *Generate Table Distribution* button.
13. After the plan has been generated successfully, and you have reviewed and are satisfied with the details, select *Execute Plan*.

    As the plan executes, you can monitor plan execution functionality.

Related Information

Save the Current Table Distribution [page 503]
Restore a Saved Table Distribution Plan [page 504]
Redistributing Tables in a Scaleout SAP HANA System

9.8.6 Table Group Advisor

Add, edit, and delete an analysis and then use the analysis in the Table Redistribution Plan Generator.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- You have the RESOURCE ADMIN system privilege.
- To preview the analysis, you also need the EXECUTE privilege on SYS.TABLE_GROUP_ADVISOR.

Context

The *Table Group Advisor* is also available within the *Table Redistribution Plan Generator* but without the preview option.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the ••• *(More)* icon on the *Table Distribution* card, and select *Table Group Advisor*.
2. (Optional) Select *Grouping Options for S4 System* to enable additional S4 specific options when adding or editing an *Analyze Existing Groups* analysis type.
3. Modify an existing analysis or click *Add Analysis* to create a new one.
   a. Specify the analysis type to use for the plan.
   b. Enter a *Preview* name.
   c. Select an analysis source.
   d. (Optional) Create an analysis filter.
   e. Select a table source.
   f. (Optional) Add tables and schemas to the analysis.
4. Click the *Preview* button to see the proposed group redistribution.
5. Click *Use in Table Redistribution.*

   The Table Redistribution Plan Generator starts. The Table Group Advisor settings are pre-populated with the values for the Table Group Analysis page. See *Generate and Execute a Table Redistribution Plan*.

Related Information

Generate and Execute a Table Redistribution Plan [page 506]

9.8.7 Table Replication

In a multi-host system, tables can be replicated on other hosts. This can reduce network traffic when, for example, slowly-changing master data often has to be joined with tables, or partitions of tables, that are located on other hosts.

For more information about table replication, see the *Table Replication* section in the *SAP HANA Administration Guide*.

9.8.7.1 Create a Table Replica

Create a replica of a partitioned or non-partitioned table on another host.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Both asynchronous table replication (ATR) and synchronous table replication (STR) are supported.

For more information on which type of replication is best suited to your needs, and to understand the limits around table replication, see the *Table Replication* section of the *SAP HANA Administration Guide*.

Procedure

1. From the *Database Overview*, click *View Current Table Distribution* from the *Table Distribution* card.

   The *Current Table Distribution* page is displayed.
2. Click the empty cell that corresponds to the table you want to replicate, and the host that you want to create the replica on.
3. In the context menu that appears, click *Create Replica Here*.
4. Choose whether you want a synchronous or asynchronous replica.

   In synchronous table replication, the source table and the its replicas always have the same state, however this results in a performance penalty when write transactions are committed.

   In asynchronous table replication, there is less of a performance penalty when committing write transactions because the source table is updated more frequently than its replicas. However, this means that data in the replica tables may be stale.

5. (For partitioned tables only) Specify the source, partitioning level, and location for the replica.

   For the source, you can choose whether to replicate the entire source table, or only specific columns.

   For the partitioning level, you can choose to use the same type of partitioning as the source table, or specify single-level partitioning.

   For the location, you can choose to automatically place the partitions, or specify their locations.
6. Choose your *Replica Schema* and *Replica Name* and click *Create*.

   The default replica schema is the source schema.

   You can choose to automatically name the replica, or create your own replica name.

9.8.7.2 Manage Table Replicas

Move, disable, or delete table replicas.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

If the replica is in the process of being synchronized, you must wait for the synchronization to be over before you can move the replica. If you want to disable or delete the replica, you can cancel the synchronization to do so.

Procedure

1. From the *Database Overview*, click *View Current Table Distribution* from the *Table Distribution* card.

   The *Current Table Distribution* page is displayed.
2. Click the table replica in the host cell that you want to manage and then click *Manage Replica*.
3. Choose one of the following options:

Option	Action
Click *Move Replica*	Choose the host to move the replica to and click *Move*. If the replica is partitioned, you must select a target host for each partition before clicking *Move*. You cannot move a replica while it is synchronizing with the source table. You must wait until the synchronization is complete.
Click *Drop Replica*	Confirm that you wan to drop the replica.
Click *Disable Replica*	Confirm that you want to disable the replica.

10 Monitoring, Analyzing, and Improving Performance

You can monitor, analyze, and improve the performance of the database using the SAP HANA cockpit.

How do you use performance monitoring tools in SAP HANA cockpit?

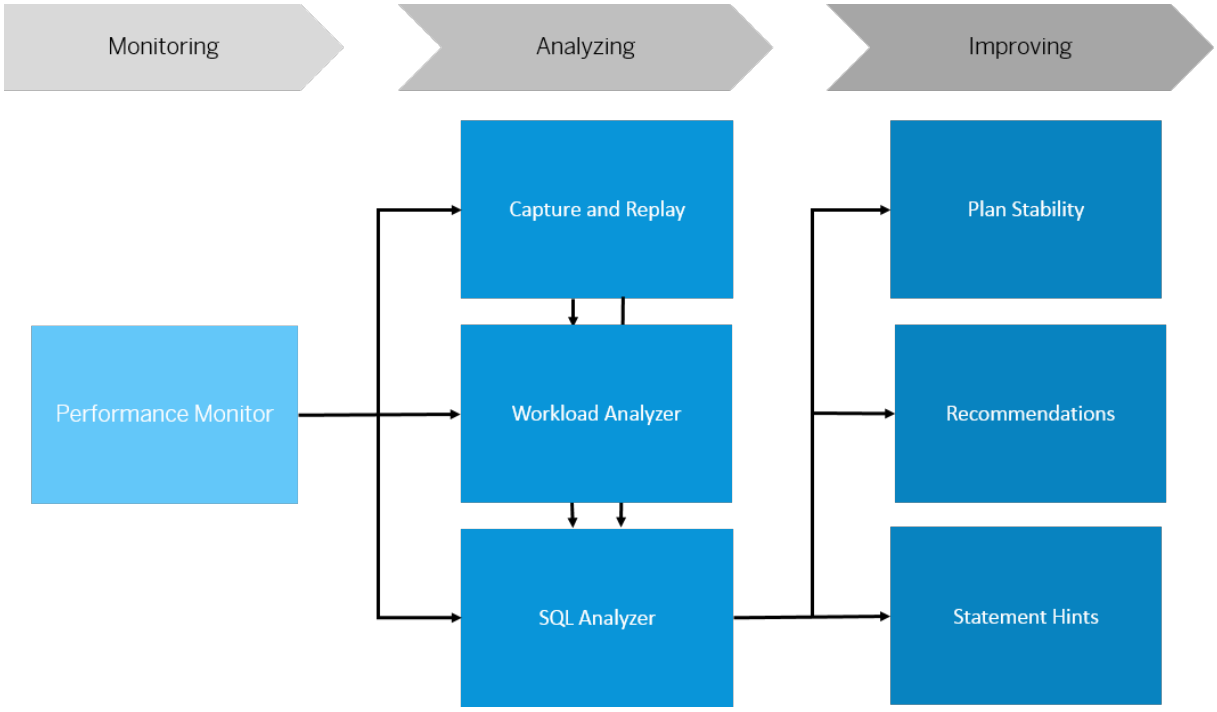

- Monitoring Performance in SAP HANA Cockpit [page 514]
- Monitoring Performance in SAP HANA Cockpit [page 514]
- Improving Performance in SAP HANA Cockpit [page 584]
- The Performance Monitor [page 515]
- Capturing and Replaying Workloads [page 532]
- Analyzing Workloads [page 561]
- Analyzing Statement Performance [page 568]
- SQL Plan Stability [page 588]
- Recommendations [page 470]
- Statement Hints [page 590]

Related Information

Monitoring Performance in SAP HANA Cockpit [page 514]
Analyzing Performance in SAP HANA Cockpit [page 531]
Improving Performance in SAP HANA Cockpit [page 584]

10.1 Monitoring Performance in SAP HANA Cockpit

Monitoring past and current information about the performance of the SAP HANA database is important for root-cause analysis and the prevention of future performance issues.

You can use the following tools to monitor fine-grained aspects of system performance in the SAP HANA cockpit:

- Use the *Performance Monitor* to visually analyze historical performance data across a range of key performance indicators related to memory, disk, and CPU usage.
- Use *Threads* to monitor the longest-running threads active in your system.You can use it to see, for example, how long a thread is running, or if a thread is blocked for a prolonged period.
- Use the *Sessions* card to monitor all sessions in your landscape.
- Use the *Statements Monitor* to analyze the current most critical statements running in the database.
- Use *Expensive Statements* to analyze individual SQL queries whose execution time was above a configured threshold.
- Use the *SQL plan cache* to get an insight into the workload of the SAP HANA database as it lists all statements currently cached in the SAP HANA database.
- Use the *Blocked Transactions* to monitor the details of transactionally blocked threads.

Related Information

The Performance Monitor [page 515]
Threads [page 521]
Session Monitoring [page 525]
Monitor and Analyze Active Statements [page 526]
Monitor and Analyze Expensive Statements [page 529]
Monitor and Analyze Statements with SQL Plan Cache [page 528]
Blocked Transactions [page 530]

10.1.1 The Performance Monitor

Analyzing the performance of the SAP HANA database over time can help you pinpoint bottlenecks, identify patterns, and forecast requirements. Use the *Performance Monitor* to visually analyze historical performance data across a range of key performance indicators related to memory, disk, and CPU usage.

What is the Performance Monitor?

The *Performance Monitor* opens displaying the load graph for the selected resource: CPU, disk, or memory. The load graph initially visualizes resource usage of all hosts and services listed on the left according to the default KPI group of the selected resource.

Change the information displayed on the load graph by:

- Switching between a number of predefined views:
  - *Default*
  - *CPU*
  - *Disk*
  - *Memory*
- Defining the monitored time frame by entering your desired dates or selecting form *Presets*.
- Setting the automatic refresh rate.
- Using the *Add Chart* button to create custom charts displaying the host and services selection, and selected KPIs. For a list of all available KPIs, see *Key Performance Indicators*.
- Filtering the results you see in the chart by selecting only the KPIs you are interested in. The hierarchically sorted table legend displays the KPI unit, y-axis scale, as well as minimum, maximum, and average values.
- Zooming into a specific time on a graph by brushing across the desired selection on the load graph directly. Click on the zoom in button on upper right corner of the highlighted area. Select *Undo* in the header toolbar to zoom out again
- Comparing the performance of your selected KPIs at different times using the *Performance Comparison* page. For more information, see *Compare Performance*.
- Clicking the ⚙ *(Settings)* icon, to customize your graphs by including hosts and services as well as additional KPIs in the *Charts* tab. In the *Alerts* tab, configure alerts according to category and priority status. You can then save the customizations as a view. See Saved Views.

How can I access the Performance Monitor?

On the *Database Overview* page, with the *Monitoring* or *All* view selected, you access the *Performance Monitor* page by clicking the title, the graph, or *Monitor Performance* on the following cards:

- *Memory Usage*
- *CPU Usage*
- *Disk Usage*

You can also click the *Monitor Performance* link on the *Monitor* card.

Related Information

Key Performance Indicators [page 519]
Compare Performance [page 516]
Performance Monitor Data for SAP Support [page 517]
Import Performance Monitor Data [page 518]
Export Performance Monitor Data [page 519]
Analyze Workloads Based on Thread Samples [page 564]
Memory Analysis [page 139]

10.1.1.1 Compare Performance

Use *Performance Comparison* to examine the performance of your selected KPIs at different time intervals.

Prerequisites

- You have the CATALOG READ system privilege.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Monitor Performance* on the *Memory Usage*, *CPU Usage*, *Disk Usage*, or *Monitoring* card.
   The *Performance Monitor* opens displaying the load graph for the selected database: CPU, disk, or memory.
2. In the *Performance Monitor*, select the KPIs you want to compare.
3. Click and drag your mouse within the graph to select a time duration.

   This selection will make up the main chart that you can contrast to any additional charts you create on the *Performance Comparison* page.
4. Click the ⚖ *(Compare)* icon in the highlighted area.
5. The *Performance Comparison* page opens, displaying the KPIs as well as hosts and services that were selected in the *Performance Monitor*.
6. Optional: You can add or remove KPIS by clicking the *Refine KPIs* button in the header toolbar and making your selection. You can also adjust hosts and services.

   Optional: You can adjust the time range of the chart by selecting the desired start and end of the monitored time interval, or choosing from *Presets* in the header toolbar.
7. Add an aditional chart for comaring performance at different time intervals by selecting the *Add a chart to compare* link on the bottom of the screen or by clicking the *Add chart* button in the header toolbar.

A selection of preset time intervals to choose from opens. Once you have made your choice, the additional chart displaying that time range appears.

> **i Note**
>
> The *Add a chart to compare* link is only available for the first additional chart, any other chart must be added by using the *Add Chart* button in the header toolbar.

8. Optional: Per default, the monitored time interval is defined via a range. To choose a time interval that is dynamically adjusted to the time interval of the main chart, click on the *Relational* button above the respective chart and make your time interval selection.
9. Optional: Update the chart by pressing the *Update* button above the respective chart.
10. Optional: You can bookmark a time range in a load chart to easily refer to it in the future.

    Highlight a time range on the desired chart, click the navigation icon on the top right corner of the highlighted area, and choose *Bookmark Selection*.

    The highlighted area changes color to indicate that a bookmark has been set. Above the chart containing bookmarks, there is a link with the number of bookmarks contained in the chart. It lists the bookmarked time range as well as the bookmark selection date. Clicking it highlights and navigates to the bookmarked time range on the chart.

    It is possible to name the bookmark by clicking the navigation icon on the highlighted area and selecting *Add Description*. The description is displayed in the bookmark list above the chart.

    To modify the description, click on the navigation icon. You can also delete the bookmark through the navigation icon or by clicking the trash bin.
11. Navigate to the *Performance Monitor* or *Workload Analysis* page by highlighting a time range, clicking on the navigation icon and making your selection.

10.1.1.2 Performance Monitor Data for SAP Support

To help SAP Support analyze and diagnose problems with your system, you can collect a snapshot of the performance monitor data from your system into a zip file. You can trigger the collection of diagnosis information from the SAP HANA cockpit.

Related Information

Import Performance Monitor Data [page 518]
Export Performance Monitor Data [page 519]

10.1.1.2.1 Import Performance Monitor Data

To analyze and diagnose problems with the SAP HANA database, you can import performance monitor data from a zip file into the SAP HANA cockpit.

Prerequisites

- You have the CATALOG READ system privilege.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Monitor Performance* on the *Memory Usage*, *CPU Usage*, *Disk Usage*, or *Monitoring* card.
   The *Performance Monitor* opens displaying the load graph for the selected database: CPU, disk, or memory.
2. Click the ↥ *(Import)* icon.

   The *Support Tools* page opens, displaying a list of available performance monitor data sets.
3. If the desired data set is listed, click the data set name to open it.
4. If the desired data set is not listed, you can import it from a local or a remote file,
   a. Select *Import* on the bottom of the page.
   b. Enter a name for the data set.
   c. Browse to the file containing the performance monitor data set that you want to import.
   d. Click *Import*.
   e. 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*.

      Click the data set name to open it.

10.1.1.2.2 Export Performance Monitor Data

To help SAP Support analyze and diagnose problems with the SAP HANA database, you can export performance monitor data into a zip file, which you can then download and, for example, attach to a support message.

Prerequisites

- You have the CATALOG READ system privilege.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Monitor Performance* on the *Memory Usage*, *CPU Usage*, *Disk Usage*, or *Monitoring* card.
   The *Performance Monitor* opens displaying the load graph for the selected database: CPU, disk, or memory.
2. Click the ⭳ *(Export All)* icon to export the CPU, disk, or memory KPI data.
3. Click *Export* in the *Export All* dialog.

   The system collects the relevant information and saves it to a zip file. This may take some time and runs in the background.

   Once the collection is available, you can download it by clicking the download button. It will be saved to the download directory of your browser on your client.

10.1.1.3 Key Performance Indicators

The *Performance Monitor* allows you select a range of host-level and service-level KPIs to analyze historical performance data of the SAP HANA database.

Host KPIs

KPI	Description
CPU	CPU used by all processes related to the operating system (OS)
Database resident memory	Physical memory used by all SAP HANA database processes

KPI	Description
Total resident memory	Physical memory used by all OS processes
Physical memory size	Total physical memory
Database used memory	Memory used by all SAP HANA database processes
Database allocation limit	Memory allocation limit for all SAP HANA database processes
Disk used	Disk space used by data, log, and trace files belonging to the SAP HANA database
Disk size	Total disk size
Network in	Bytes read from the network by all processes
Network out	Bytes written to the network by all processes
Swap in	Bytes read from swap memory by all processes
Swap out	Bytes written to swap memory by all processes

Services KPIs

KPI	Description
CPU	CPU used by the database process
System CPU	CPU used by the database process relative to the operating system
Memory used	Memory used by the database process
Memory allocation limit	Effective allocation limit of the database process
Handles	Number of open handles in the index server process
Ping time	Indexserver ping time including nsWatchdog request and collection of service-specific KPIs
Swap in	Bytes read from swap by the process
Open connections	Number of open SQL connections
Open transactions	Number of open SQL transactions
Blocked transactions	Number of blocked SQL transactions
Statements	Number of finished SQL statements
Active commit ID range	Range between newest and oldest active commit ID
Pending session request count	Number of pending requests
Active versions	Number of active MVCC versions
Acquired record locks	Number of acquire record locks
Read requests	Number of read requests (selects)
Write requests	Number of write requests (insert, update, and delete)
Merge requests	Number of merge requests
Column unloads	Number of table and column unloads

KPI	Description
Active threads	Number of active threads
Waiting threads	Number of waiting threads
Total threads	Total number of threads
Active SqlExecutors	Total number of active SqlExecutor threads
Waiting SqlExecutors	Total number of waiting SqlExecutor threads
Total SqlExecutors	Total number of SqlExecutor threads
Data write size	Bytes written to data area
Data write time	Time used for writing to data area
Log write size	Bytes written to log area
Log write time	Time used for writing to log area
Data read size	Bytes read from data area
Data read time	Time used for reading from data area
Log read size	Bytes read from log area
Log read time	Time used for reading from log area
Data backup write size	Bytes written to data backup
Data backup write time	Time used for writing to data backup
Log backup write size	Bytes written to log backup
Log backup write time	Time used for writing to log backup
Mutex Collisions	Number of collisions on mutexes
Read/Write Lock Collisions	Number of collisions on read/write locks

Related Information

Memory Usage in the SAP HANA Database

10.1.2 Threads

Use *Threads* to monitor the longest-running threads active in your system. You can use it to see, for example, how long a thread is running, or if a thread is blocked for a prolonged period.

Analyzing the threads running in the SAP HANA database can be helpful when analyzing the current system load.

You can identify which statements or procedures are being executed and at what stage they are, who else is connected to the system, and if there are any internal processes running as well.

The *Threads* card provides information about the number of currently active and blocked threads in the database.

On the *Database Overview* page, with the *Monitoring* or *All* view selected, on the *Threads* card, click either the number of active or blocked threads on the card or the card title to open the *Threads* page.

The *Threads* page allows you to monitor the longest-running threads in your current system. You can retrieve more information or customize what is being displayed, for example:

- Filter threads by host, service, and thread type
- Choose the sorting order by checking the *Group and Sort* box and selecting the sorting parameters
- See the call stack information on your chosen thread
- Define columns and choose the parameters you want information on

When you have selected a thread, you can *Navigate To* the *Sessions* or *Blocked Transactions* page for the thread with the same connection ID.

If a thread is in a blocked transaction or is using an excessive amount of memory, you can cancel the operation executing the thread by clicking *Cancel Operations* in the footer toolbar.

Related Information

Thread Details [page 522]

10.1.2.1 Thread Details

The *Threads* card 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 Columns dialog, which you open by clicking the ⚙ *(Settings)* icon in the table toolbar.

Thread Information

To open the *Threads* card, on the *Database Overview* page, with the *Monitoring* or *All* view selected, click either the number of active or blocked threads on the card or the card title.

The table below lists the information available for threads.

Detail	Description
Blocking Transaction	Blocking transaction
Duration (ms)	Duration (ms)
Host	Host name

Detail	Description
Port	Internal port
Service	Service name
Hierarchy	Thread grouping information. Filled with Connection ID/ Update Transaction ID/Transaction ID or left empty for inactive threads
Connection ID	Connection ID
Thread ID	Thread ID
Calling	The thread or service which the thread calls
Caller	The thread or service which called this thread
Thread Type	Thread type
Thread Method	Thread method
Thread Detail	Thread detail
User	User
Application User	Application user name
CPU Time	CPU time of thread
Cumulative CPU Time	CPU time of thread and associated children
Transaction ID	Transaction ID
Update Transaction ID	Update transaction ID
Thread Status	Thread state
Connection Transaction ID	Transaction object ID
Connection Start Time	Connected Time
Connection Idle Time (ms)	Time that the connection is unused and idle
Connection Status	Connection Status: 'RUNNING' or 'IDLE'
Client Host	Host name of client machine
Client IP	IP of client machine
Client PID	Client Process ID
Connection Type	Connection type: Remote, Local, History (remote), History (local)
Own Connection	Own connection: TRUE if own connection, FALSE if not
Memory Size per Connection	Allocated memory size per connection
Auto Commit	Commit mode of the current transaction: TRUE if the current connection is in auto-commit mode, FALSE otherwise

Detail	Description
Last Action	The last action done by the current connection: ExecuteGroup, CommitTrans, AbortTrans, PrepareStatement, CloseStatement, ExecutePrepared, ExecuteStatement, FetchCursor, CloseCursor, LobGetPiece, LogPutPiece, LobFind, Authenticate, Connect, Disconnect, ExecQidItab, CursorFetchItab, InsertIncompleteItab, AbapStream, TxStartXA, TxJoinXA
Current Statement ID	Current statement ID
Current Operator Name	Current operator name
Fetched 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, TrexViaDbsl from ABAP, HybridTable Reorganizer, Session external
Is Encrypted	Encrypted: TRUE if the secure communication is enabled (SSL enabled), FALSE, otherwise
Connection End Time	The time when the connection is closed for history connections
Blocked Update Transaction ID	Write transaction ID of the write transaction waiting for the lock
Blocking Transaction ID	Transaction object ID of the transaction holding the lock
Thread ID of Lock Owner	Connection ID associated with the blocked write transaction
Blocking Update Transaction ID	Write transaction ID of the write transaction holding the lock
Transactional Lock Type	Transactional lock type
Transactional Lock Mode	Transactional lock mode
Lock Wait Component	Waiting for lock component
Lock Wait Name	Waiting for lock ID
Timestamp of Blocked Transaction	Timestamp of the blocked transaction
Waiting Record ID	ID of the record on which the lock is currently placed
Waiting Object Name	Name of the object on which the lock is currently placed
Waiting Object Type	Type of the object on which the lock is currently placed

Detail	Description
Waiting Schema Name	Name of the schema on which the lock is currently placed

10.1.3 Session Monitoring

Use the *Sessions* card to monitor all sessions in your landscape.

Analyzing the sessions connected to your SAP HANA database helps you identify which applications or which users are currently connected to your system, as well as what they are doing in terms of SQL execution.

The *Sessions* card displays the number of active and total sessions.

On the *Database Overview* page, withthe *Monitoring* or *All* view selected, on the *Sessions* card, click either the number of active or blocked threads on the card or the title of the card to open the *Sessions* page.

The *Sessions* page allows you to monitor all sessions in the current landscape. You can see the following information:

- Active/inactive sessions and their relation to applications
- Whether a session is blocked and, if so, which session is blocking it
- The number of transactions that are blocked by a blocking session
- Statistics such as average query runtime and the number of DML and DDL statements in a session
- The operator currently being processed by an active session

To support monitoring and analysis, you can perform the following actions on the *Sessions* page:

- Cancel a session by choosing *Cancel Sessions*
- Save the data sets as a text or html file by choosing the *Save As...* button.

10.1.4 Statements

Use *Monitor Statements* to monitor and analyze different types of statements in your system.

On the *Database Overview*, with the *Monitoring* or *All* view selected, click the *SQL Statements* card to open the *Monitor Statements* page.

There are four views available in the table on the *Monitor Statements* page:

Monitor Statements Table

View	Function
Overview	An overview and analysis of the most critical statements running in the database.
Active Statements	A list of all statements currently running in the system.
SQL Plan Cache	Insights into the workload of the SAP HANA database through a list of all statements currently cached in the SAP HANA database

View	Function
Expensive Statement Trace	Analysis of individual SQL queries whose execution time is above a configured threshold.

You can use the search bar above the table to search for and display only the SQL statements you are interested in.

Related Information

Monitor and Analyze Active Statements [page 526]
Monitor and Analyze Statements with SQL Plan Cache [page 528]
Monitor and Analyze Expensive Statements [page 529]
Session Monitoring [page 525]

10.1.4.1 Monitor and Analyze Active Statements

On the *SQL Statements* page, use the *Overview* view to analyze the current most critical statements running in the database, and the *Active Statements* view to see all actively running statements.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Analyzing the current most critical statements running in the SAP HANA database can help you identify the root cause of poor performance, CPU bottlenecks, or out-of-memory situations. Enabling memory tracking (*Overview* view only) allows you to monitor the amount of memory used by single statement executions.

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. On the *Database Overview* page, with the *Monitoring* or *All* view open, click the title or *View all* on the *SQL Statements* card.

   The *Overview* view opens. It allows you to analyze the most current statements running in the database. You can see:

   - The 100 most critical statements, listed in order of the longest runtime
   - The full statement string and ID of the session in which the statement is running
   - Application, application user, and the database running the statement
   - Whether a statement is related to a blocking transaction

2. Click the *Active Statements* tab.

   The *Active Statements* view shows:

   - The full statement string and ID of the session in which the statement is running
   - The host and port, as well as the connection ID of the active statement
   - The last executed time
   - If the statement has a parent statement and its ID

3. To support monitoring, you can perform these actions on either view:
   a. See the full SQL statement and open the statement with the *SQL analyzer* by clicking *More* in the *Statement String* column and selecting *Open in SQL Analyzer*.
   b. Navigate to the *Sessions* page to analyze the session with the same ID as the statement by clicking on the *(Navigate to Sessions)* icon next to the workload class column.
   c. Set up or modify workload classes by clicking a statement's *Workload Class Name*. Choose *New* to create a new workload class or *Existing* to select a workload class from a list, then fill in the fields.
   d. (*Overview* view only) Enable or disable memory tracking.
   e. Personalize the columns displayed in the view using the *(Settings)* icon.

Related Information

Setting a Memory Limit for SQL Statements
Create a Workload Class [page 460]
Create a Workload Class Mapping [page 462]

10.1.4.2 Monitor and Analyze Statements with SQL Plan Cache

On the *SQL Statements* page, use the *SQL Plan Cache* view to get an insight into the workload of the database as it lists all statements currently cached in the SAP HANA database.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Analyzing all statements currently cached in the SAP HANA database can help you identify statement hashes, as well as if a statement has been correctly cached.

Technically, the plan cache stores compiled execution plans of SQL statements for reuse, which gives a performance advantage over recompilation at each invocation. For monitoring reasons, the plan cache keeps statistics about each plan, for instance number of executions, min/max/total/average runtime, and lock/wait statistics. Analyzing the plan cache is very helpful as one of the first steps in performance analysis because it gives an overview about what statements are executed in the system.

> **i** Note
>
> Due to the nature of a cache, seldom-used entries are evicted from the plan cache.

The SQL plan cache is useful for observing overall SQL performance as it provides statistics on compiled queries. You can get insight into frequently executed queries and slow queries with a view to finding potential candidates for optimization.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view open, click the title or *View all* on the *SQL Statements* card to open the *SQL Statements* page.
2. Select the *SQL Plan Cache* tab.
3. To support monitoring and analysis, you can perform the following actions on the view:
   a. Open the selected SQL statement with the SQL analyzer by clicking *More* next to the statement string and selecting Open in SQL Analyzer.
   b. Open the selected SQL statement with SQL analyzer and save the executed plan statistics as a PLV file by clicking on *More* next to the statement string and selecting *Analyze and Save Plan*.
   c. Save the data sets as a text or HTML file by choosing the *Save As...* button.

d. Configure the plan cache size by choosing *Configure*.
e. Clear all cached statements by choosing t *Clear All Plan Cache*.
f. Personalize the columns displayed in the view using the ⚙ *(Settings)* icon.

10.1.4.3 Monitor and Analyze Expensive Statements

On the *SQL Statements* page, use the *Expensive Statements* view to analyze individual SQL queries whose execution time is above a configured threshold.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Analyzing expensive statements can help you understand why they exceed duration thresholds.

The expensive statements trace records information about the expensive statements for further analysis and displays it on the *Expensive Statements* page.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click the *SQL Statements* and then click the *Expensive Statements* tab.
2. To support monitoring and analysis,you can perform these actions on the view:
   a. Define the monitored date.
   b. Filter expensive statements, refresh the list, choose the sorting parameter, and filter by parameter.
   c. Save the data sets as a text or HTML file by choosing the *Save As...* button.
   d. Open an expensive statement with the SQL analyzer by clicking *More* in the *Statement String* and selecting *Open in SQL Analyzer*.
   e. Set up or modify workload classes by clicking a statement's *Workload Class Name*. Choose *New* to create a new workload class or *Existing* to select a workload class from a list, then fill in the fields.
   f. Access the *Current Table Distribution* page to view the current table distribution in detail.
   g. Personalize the columns displayed in the view using the ⚙ *(Settings)* icon.

Related Information

Expensive Statements Trace
Create a Workload Class [page 460]
Create a Workload Class Mapping [page 462]

10.1.5 Blocked Transactions

Use *Blocked Transactions* to monitor transactionally blocked threads. You can use it to see, for example, what transaction is blocking a thread, the type of lock held, and the owner of the lock.

Blocked transactions are transactions that are unable to be processed further because they need to acquire transactional locks (record or table locks) that are currently held by another transaction. Transactions can also be blocked while waiting for other resources such as network or a disk (database or metadata locks).

Analyzing the blocked transactions in SAP HANA database can be helpful when analyzing the current system load, as transactionally blocked threads can impact application responsiveness.

The *Blocked Transactions* feature provides information on the number of currently blocked threads in the database.

On the *Database Overview*page, with the *Monitor* or *All* view selected, click the *Open Blocked Transactions* link on the *Monitoring* card.

To support monitoring and analysis, you can perform the following actions on the *Blocked Transactions* page:

- Filter transactions with the help user-defined keywords
- Select to hide own or idle sessions
- Customize the blocked transaction columns to show only desired parameters
- Click on a blocked transaction and select *Navigate To...* on the bottom right of the screen to jump to *Threads* or *Sessions* with the same connection ID.

10.1.6 Monitor Table Usage

Monitor tables to optimize resource utilization and improve query performance.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

With *Table Usage* you can visualize tables by size, explore the usage history of tables, and move tables to warm storage.

Procedure

1. On the *Database Overview* page, with the *Monitoring* or *All* view selected, click the *Monitor table usage* link on the *Monitoring* card.
2. To filter tables shown, adjust the *Total Access*/*Size*/*Display* values and click *Go*. Click *Reset* to remove filters.

   For the best display, select up to 50 tables. Two options for table analysis are available:

   - For table format display, choose *(Table Chart)*.
   - For graphical format display, choose *(Bubble Chart)* . 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.

Related Information

SAP Note 2092669
Alerts [page 126]

10.2 Analyzing Performance in SAP HANA Cockpit

You can analyze the performance of the database using the SAP HANA cockpit.

You can use the following tools to analyze fine-grained aspects of database performance in the SAP HANA cockpit:

- Use the capture and replay to detect, analyze, or verify any potential issues before applying changes or upgrades.

- Use the workload analyzer to analyze the database performance through data from thread samples. You can also use it to analyze the workload captured with the capture and replay tool, or any other workload occurring in a system.
- Use the SQL analyzer to understand performance issues of a query execution and other query execution aspects of the SAP HANA database.

Related Information

Capturing and Replaying Workloads [page 532]
Analyzing Workloads [page 561]
Analyzing Statement Performance [page 568]

10.2.1 Capturing and Replaying Workloads

Capturing and replaying workloads from an SAP HANA database helps you evaluate potential impacts on performance or stability after a change in the hardware or software configuration.

The following sections provide an overview of SAP HANA capture and replay:

What can I do with SAP HANA capture and replay?

This tool allows you to capture the workload of a source system and to replay the captured workload on a target system without applications.

Moreover, you can use the tool to analyze the captured workload and the reports generated after replaying the workload. Comparing the performance between the source and target systems can help you find the root cause of performance differences.

What is a workload?

Workload in the context of SAP HANA can be described as a set of requests with common characteristics. For more information about workload, see *Workload in the Context of SAP HANA*.

In the context of SAP HANA capture and replay, workload can mean any change to the database via SQL statements that come from SAP HANA client interfaces such as JDBC, ODBC, or DBSL. The workload can be created by applications or clients (for example, SAP NetWeaver or Analytic).

When should I use SAP HANA capture and replay?

Use SAP HANA capture and replay to detect, analyze, or verify any potential issues before applying changes or upgrades, such as:

- Hardware change
- SAP HANA revision upgrade
- SAP HANA ini file change
- Table partitioning change
- Index change
- Landscape reorganization for SAP HANA scale-out systems
- Apply HINT to queries

For more information on possible use cases, see *An overview of possible use cases for SAP HANA capture and replay*.

How can I access SAP HANA capture and replay?

Open SAP HANA capture and replay from the SAP HANA cockpit. On the *Database Overview* page, search for the *Capture Workload* and *Replay Workload* tiles.

How does SAP HANA capture and replay work?

The main steps involved in the capturing and replaying process are:

1. Capture
   In this step the tool automatically collects the execution context information together with the incoming requests to the database. The captured workload file stores the start times of the SQL statements.
   A database backup is recommended after starting capturing to ensure that the source and target systems are in a consistent state.
2. Preprocess
   In this step the tool reconstructs and optimizes the captured workload file to make it replayable on a target system. This process is a one-time operation and the stored preprocessed workload file can be replayed multiple times.
3. Replay
   The replayer is a service on operating system level that needs to be started before replaying.
   The tool replays the preprocessed file based on the SQL statement timestamp or on the transactional order. Together with the collected execution context it allows you to accurately simulate the database workload.
4. Analyze
   For a final analysis, you can generate comparison reports displaying a capture-replay or a replay-replay comparison. You can analyze the statements based on results or on performance.

Related Information

Landscape Considerations [page 534]
System Privileges [page 538]
Files and Sizing Guidelines [page 539]
Capturing a Workload [page 541]
Replaying a Workload [page 546]
Workload in the Context of SAP HANA
Security Risks of Trace, Dump, and Captured Workload Files
https://blogs.sap.com/2019/01/15/an-overview-of-possible-use-cases-for-sap-hana-capture-and-replay/
SAP Note 2362820

10.2.1.1 Landscape Considerations

A 3-tier setup is recommended.

The 3-tier setup has the following advantages:

- Three or more servers are possible depending on the number of dedicated replayer systems.
- The replay results will be stored in the separate control system when recovering the target system.
- This setup is recommended for replays with re-initialization of the target system.
- This setup is recommended for scenarios with multiple replayers.

The graphic below offers an example of a 3-tier setup. Keep in mind that the replayers and SAP HANA cockpit can run on the control system server or on individual servers.

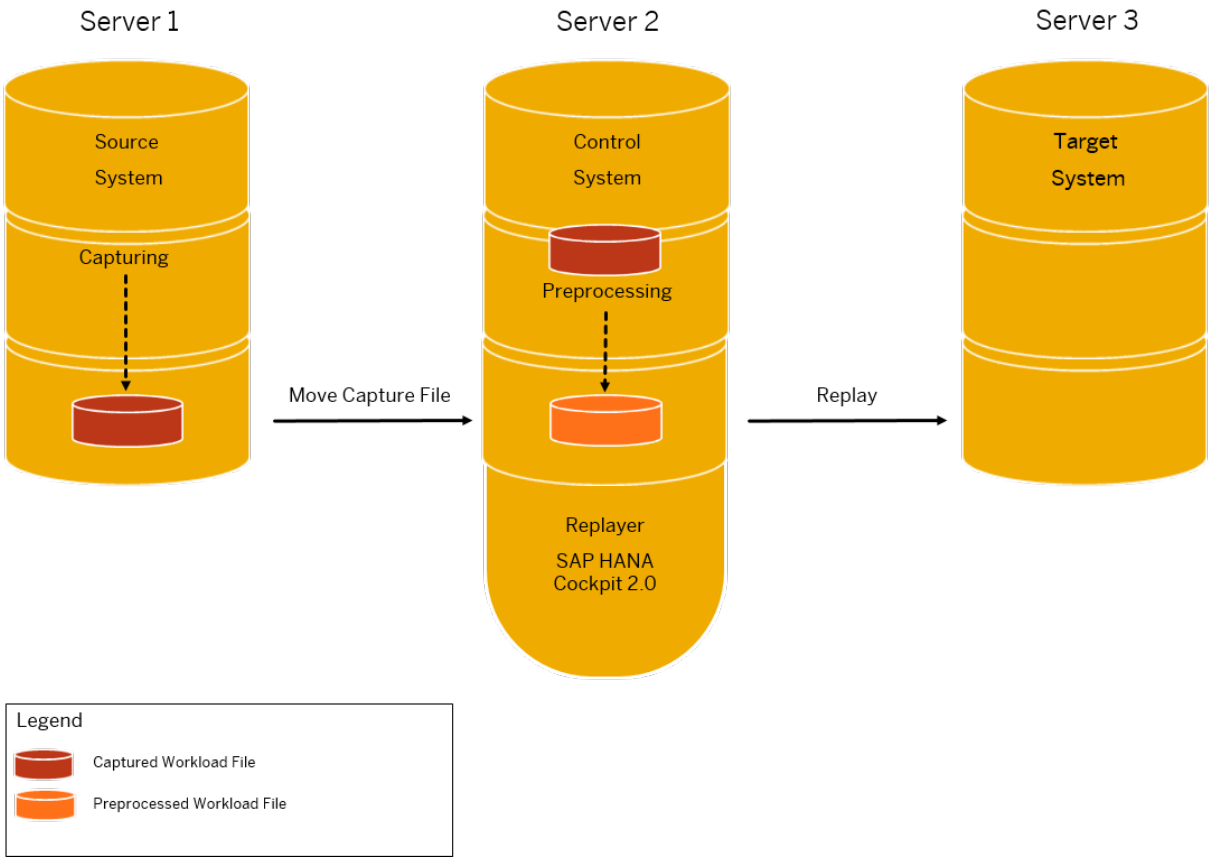

> **i** Note
>
> SAP HANA capture and replay can also be used with SAP HANA, Express Edition installations. For more information, see *Adjust the Global Allocation Limit* in the *Getting Started with SAP HANA 2.0, express edition (Virtual Machine Method)* or *Capture and replay with multiple distributed replayers on SAP HANA, Express Edition*.

System types

These systems should be part of a 3-tier setup:

System Type	Description
Source System	• The original workload is captured here • The captured workload file and the database backup are initially created here and stored on disk **i Note** Running a capture will have an overhead on CPU utilization and disk I/O.
Control System	• This can be a dedicated system or a shared system • Manages and directs replayers when running a replay • Stores replay results for analysis • Preprocessing is executed here • Serves as central monitoring system for the replay process **i Note** Preprocessing and replaying require CPU and memory to execute.
Target System	• This can be a dedicated system or a shared system • The replay is executed here, even though it's started from the control system • For meaningful replay results, the system should be in consistent state with the source system (recovered via database backup from the source system) **i Note** If the control system and target system are the same, the replay results are lost when the target system is recovered.
Replayer	• Process running on the operating system • Is responsible for reading from preprocessed workload files and executing workload on the target system • Replayers can run on individual systems • Multiple replayers can be used to scale the workload during replay • Replayers can run on SAP HANA, express edition installations **→ Recommendation** Run the replayers on the control system or a separate system because the replay requires CPU and memory to execute.

System Type	Description
SAP HANA cockpit 2.0	• Used to set up, configure, monitor, analyze while capturing, preprocessing and replaying • SAP HANA cockpit 2.0 can run on SAP HANA, express edition installations **→ Recommendation** Run SAP HANA cockpit 2.0 on the control system host.

> **i Note**
>
> The replayers and SAP HANA cockpit can run on the control system server or on individual servers.

Prerequisites

- Check the disk performance to ensure that there is sufficient bandwidth for capturing and preprocessing workloads without any performance bottlenecks. If disk performance is not sufficient, the active capture can impact the source system.
- Check the available disk space in combination with the characteristics of the workload that should be captured. The required disk space is highly dependent on the type of workload being captured.
  Use the disk space that is dedicated to the database instance itself.
- One replayer service is sufficient to execute a replay successfully. For a better scalability and performance in large workload scenarios, multiple replayers can be used for all replaying purposes. When using multiple replayers, distribute and divide all involved components (for example, target instance, control instance, one or more replayers) on different hosts and systems. Doing so will reflect the initial captured workload as realistic as possible. This will also reduce the effect which the resource consumption of the components might have on a replay.
- Use a separate control and target instance for replaying workloads. If a replayed statement causes a crash, it will be displayed in the replay report. When you use one and the same control and target instance, the replay report entry causing a crash will not be successfully sent to the control instance.
- Use the secure store for saving passwords and authenticating users. For more information, see *Secure User Store (hdbuserstore)* in the *SAP HANA Security Guide*.
- The target system should meet the same privacy and security prerequisites as the source system. Since the target system processes the same data as the source system, it should meet an appropriate security level depending on data criticality.
  Unnecessary network connections to the target system should not be allowed. Users registered on the source system might be able to access the target system after a replay has been completed.
- Regarding version dependencies, the following rules can be followed:
  - Target system >= Control system & Replayers >= Source system
  - The source system should be at least 122.14+ for captures with transactional replay enabled
- To trigger replays, the control system and target system must be registered in the same SAP HANA cockpit. The user in the SAP HANA cockpit should be able to access them both.
  When registering the target system, the cockpit does not store the credentials. For more information, see *Register a Database*.

For more information, see *SAP Note 2362820*.

Related Information

Secure User Store (hdbuserstore)
https://blogs.sap.com/2017/12/20/capture-and-replay-with-multiple-distributed-replayers-on-sap-hana-express-edition/
SAP Note 2362820
SAP HANA Security Guide
Registered Databases [page 47]

10.2.1.2 System Privileges

You need the following system privileges when using SAP HANA capture and replay.

System and User	Description
Cockpit User for SAP HANA cockpit	• Requires the Cockpit Database Administrator role • It is used for registering, managing and accessing databases
Workload Capture Admin for the source system	• Used to start captures and trigger backups in the source system • Requires the WORKLOAD CAPTURE ADMIN privilege to capture workloads • Optional: BACKUP OPERATOR privilege to trigger backups • Optional: INIFILE ADMIN privilege to see the previously used optional filters on the capture configuration page • This user needs to be used as the connecting user when registering the source system in the same SAP HANA cockpit
Control Replay Admin for the control system	• Used for preprocessing and replaying in the control system • Requires the WORKLOAD REPLAY ADMIN privilege for preprocessing and replaying workloads, as well as generating replay reports • Requires the WORKLOAD ANALYZE ADMIN privilege for loading and analyzing workloads • Requires the CATALOG READ privilege for generating replay reports • This user needs to be used as the connecting user when registering the control system in the same SAP HANA cockpit
Target Replay Admin for the target system	• Used to execute replays and reset user passwords in the target system • Requires the WORKLOAD REPLAY ADMIN privilege to execute replays • This user needs to be used as the connecting user when registering the target system in the same SAP HANA cockpit

10.2.1.3 Files and Sizing Guidelines

The files used in the process of capturing, preprocessing and replaying a workload can have different sizes depending on various factors.

Captured workload file

After capturing the workload from a source system, a captured workload file will be available for the replay. This captured workload file containes multiple captured workload segment files as indicated in the graphic below:

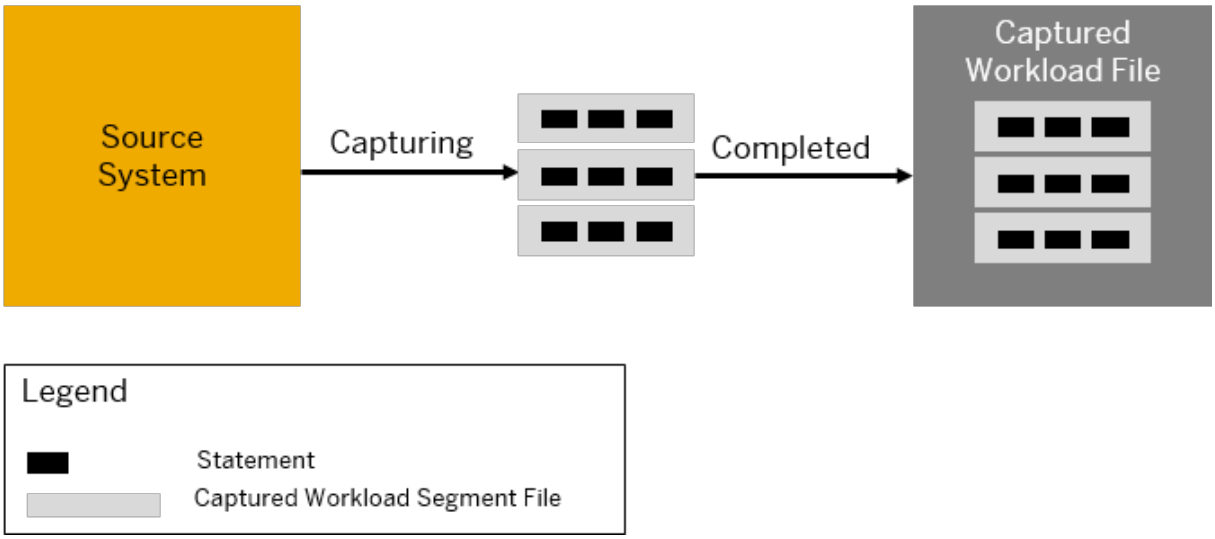

On a conceptual level, we will refer to the captured workload file using the shorter term "captured workload".

The size of the captured workload file depends on:

- Number of requests
- String size of captured statements
- Number and size of captured input parameters for statements
- If configured, the average size of the captured explain plans also needs to be considered

Activating a capture may cause:

- CPU overhead on the source system
- Memory overhead on the source system
  To configure the memory used, see the `compressionbuffersize` and `filebuffersize` parameters in *SAP Note 2362820*
- Increase in disk I/O
  The amount depends on the data that needs to be written to disk.

Preprocessed workload file

To optimize the captured workload file before replaying it, you need to preprocess it first.

While preprocessing the captured workload file, the captured statement segment files are stored in a directory. After the preprocessing is completed, the output is a preprocessed workload file containing the directory with multiple files as indicated in the graphic below:

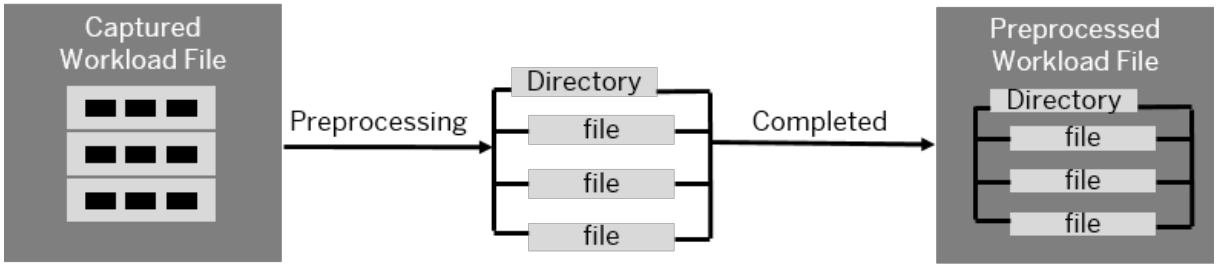

On a conceptual level, we will refer to the preprocessed workload file using the shorter term "preprocessed workload".

The size of the preprocessed workload file is larger than the captured workload file. The size varies depending on the compression and content of the captured workload file. Preprocessed workload files are always uncompressed.

Preprocessing a captured workload file may cause:

- CPU utilization
  The database will utilize the CPU as much as possible to achieve the best performance.
- Memory consumption
- Increase in disk I/O
  The amount depends on the data that needs to be read and written.

Replay

The replay process is performed by the replayer, which should be running before starting the replay as indicated in the graphic below:

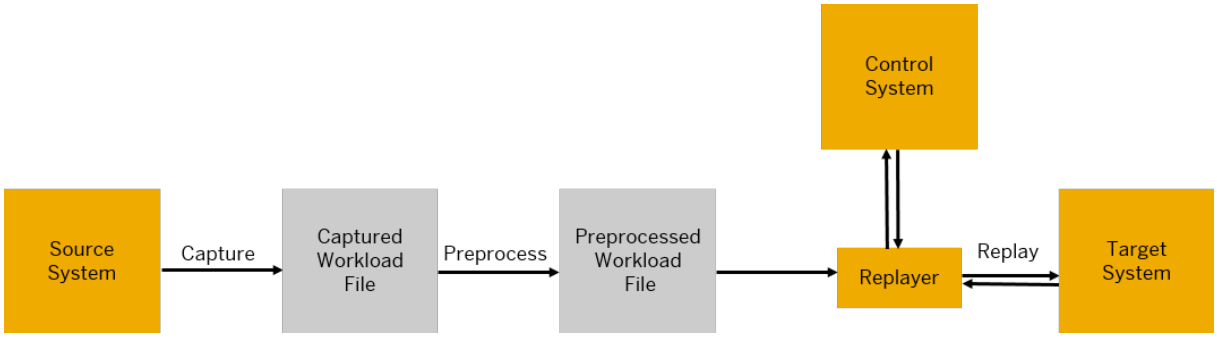

Running a replayer may cause:

- CPU utilization
- Memory consumption
  The amount of memory consumed depends on the size of the preprocessed workload file. However, even though the replayer is a separate process running on the operating system of a SAP HANA installation, its memory can be limited using the global allocation limit of its host database.
- Increase in disk I/O
  The amount depends on the data that needs to be read from disk.

The duration of the replay depends on the duration of the capture.

Running a replay may cause:

- CPU utilization
  The target system will use resources according to the workload being replayed.
- Memory consumption
  The amount of memory consumed depends on the content of the preprocessed workload file.
- Increase in disk I/O on the control system
  The amount depends on the data that needs to be written and read and is influenced by the replay result size.

The size of the replay result depends on:

- The total record count
- The memory allocated for each record
- If configured, the average size of the captured explain plans also needs to be considered

> **i Note**
>
> When using different servers, capturing, preprocessing, and replaying may cause network traffic.

Related Information

SAP Note 2362820

10.2.1.4 Capturing a Workload

You can capture the entire workload from a source system or only a part of this workload.

To capture the workload from a source system, use the *Capture Workload* card.

Related Information

Capture a Workload [page 542]
Monitoring a Captured Workload [page 545]
Capturing Configuration Settings [page 543]

10.2.1.4.1 Capture a Workload

You can capture the workload from a source system.

Prerequisites

You have the Workload Capture Admin user for the source system.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, you can choose the *Capture Workload* card or choose *Start New Capture* directly from the card.

   If you choose the *Capture Workload* card, the *Capture Management* page opens. If you already captured workload with SAP HANA capture and replay, you see the captured workload located in the currently configured capture destination. You can also see information such as name, status, start time, size, duration, or capture usage.

   If you choose *Start New Capture* directly from the card, the *Configure New Capture* page opens.
2. Optional: To change the capture destination, choose *Configure Capture* on the bottom right of the *Capture Management* page.

   The captured workload file is stored by default in the $SAP_RETRIEVAL_PATH/trace directory with a *.cpt file extension. For tenant databases there is a subfolder in the trace directory for each tenant. Since the default trace directory generally resides in the same storage area with data and log volumes, capturing workloads may affect the performance across the entire system over time. If you enter a different destination for the captured workload file, you may have a better distribution of the disk I/O between the data and log volumes and the captured workload file.
3. To start configuring the new capture, choose *New Capture* on the *Capture Management* page on the bottom right.

   On the *Configure New Capture* page it is mandatory to enter the name of the new capture.

   You can customize other optional settings before you start the capture. For an overview of these settings, see *Capture Configuration Settings*.
4. Choose *Start Capture* on the bottom right.

   The new captured workload is displayed on the *Capture Management* page with the status *Capturing*. You can stop the capture or you can let it run as long as you wish.

Related Information

Capturing Configuration Settings [page 543]
System Privileges [page 538]
Monitoring a Captured Workload [page 545]

10.2.1.4.1.1 Capturing Configuration Settings

You can customize several optional settings before you start capturing a workload.

The tables provide an overview of the settings that can be customized.

General Information

Setting	Description
Usage	Select if you will use the captured workload for a replay, an analysis, or both. Your selection in this field impacts the collected data. For example, when capturing a workload for analysis, the *Workload Details* option is automatically turned on in the *Data Collection* section.
Capture Name	Enter a name for the capture. This field is mandatory.
Description	Enter a description of the capture for future reference.
Schedule	You can schedule the capture to start at a specific point in time by specifying the start and end time. If you turn it on, the status of your captured workload changes to *Scheduled*. After setting a specific time, you can change it later. However, once you schedule a capture, you can't turn off the schedule setting anymore.

Capture Control

Setting	Description
Create Backup	Turn it on to automatically create a database backup of the source system after starting the capture. → Recommendation To ensure that the source system and the target system are in a consistent state for capture and replay, we recommend to perform a database backup after starting the capture. A database backup is required only for the first time, because incremental backups can be used once the system has been initialized for the first time. For more information about backups, see *SAP HANA Backup and Recovery* in the *SAP HANA Administration Guide*. Use the *Backup Settings* link to choose the backup type (for example, complete, differential, incremental), to select between Backint integration or file-based backups, to add a prefix to the file name, to define a path to store the backup file, or to enter parameters for the Backint. By creating a backup, you can ensure the ability to use the *Synchronize with backup* option during replay. For more information, see *Replay Configuration Settings*.
Overwrite Capture When Time Exceeds	Turn it on and enter a time to remove the captured workload segment files that are older than the specified time you entered. Only closed segments are deleted. The currently active captured workload segment file is not affected.
Overwrite Capture When Disk Usage Exceeds	Turn it on and select a ratio to remove the old captured workload segment files when the disk usage exceeds the specified percentage. Only closed segments are deleted. The currently active captured workload segment file is not affected.

Data Collection

Setting	Description
Explain Plan	Turn it on to collect the output of the EXPLAIN PLAN command for the captured statements. You can use this information for analysis after the replay.
Workload Details	Turn it on to collect additional information for the workload analyzer such as application source, involved threads, network statistics, or related objects. If this option is disabled, the captured workload file can still be viewed using the workload analyzer, but less information will be available for the review.
SQL Input Parameters	Turn it on to see the parameter values in the replay report. To do so, turn on the *Workload Details* option first. This allows you to enable the *SQL Input Parameters* option. This option is needed for replaying.

Optional Filter

Setting	Description
Optional Filter	Use the additional filters to capture only desired aspects of the workload. Filters can include different aspects such as: • *Application Name* (for example, HDBStudio) • *Application User Name* is the name of the user logged in to the application. • *Database User Name* is the name of the database user (for example, SYSTEM). • *Client* is the ABAP client number (for example, 000).

Related Information

SAP HANA Database Backup and Recovery
Replaying Configuration Settings [page 551]

10.2.1.4.2 Monitoring a Captured Workload

You can monitor a captured workload.

To view monitoring information such as duration, the number of captured statements, or disk space, open the *Capture Monitor* by choosing the started capture from the *Capture Management* page. If you defined any filters, these can also be viewed in the *Capture Information* section.

If you didn't create a backup when starting the capture, you can also start a backup from the *Capture Monitor* page. If you do so, you can check any details about the backup by clicking the drop-down arrow next to *Capture Monitor* and opening the previously started backup from the *Related Apps* section.

Back on the *Capture Management* page, you can see the capture's name, status, duration, number of statements, size, or usage. You can filter captured workloads by the start time.

You can use the captured workload file for replay or analyzing. From the *Capture Management* page, you can open the workload analyzer. It is mandatory to load the captured workload before opening it with the workload analyzer. You can do this in two ways:

- You can load the captured workload by clicking *Start* in the *Workload Analysis* column. Then choose the captured workload to open the *Workload Analysis* page.
- Alternatively, you can choose a captured workload which was not loaded before. Load it after opening the *Workload Analysis* page using the *Load* button on the top right.

For more information, see *Analyze Workloads Based on Captured Workloads*.

You can monitor the captured workload in the M_WORKLOAD_CAPTURE system view. For more information, see *M_WORKLOAD_CAPTURES System View* in the *SAP HANA SQL Reference Guide* guide.

Related Information

Analyze Workloads Based on Captured Workloads [page 566]
M_WORKLOAD_CAPTURES System View

10.2.1.5 Replaying a Workload

You can replay the preprocessed workload based on the SQL statement timestamp or on the transactional order.

Replaying a workload implies that the captured statements are executed again.

→ Recommendation

Manually copy the captured workload files from the source system to the control system and the database backup from the source system to target system.

You can replay all captured workloads as often as necessary.

→ Recommendation

When running consecutive replays, restore the target system back to a consistent state after a replay and before running another replay. This is necessary because after replaying a workload on a system, any changes applied during that replay will remain active in the system.

Example

Let's assume the captured workload file includes the statement `<INSERT INTO TABLE A VALUES (x).` During a replay, value `x` will be inserted into table A. At the end of the replay, table A contains value `x`. If you run another replay without resetting the system to its initial state, table A will contain duplicate values `x,x` at the end of the replay or the statement will fail (for example, in the case of unique constraint errors). As `x,x` does not reflect the intended end result of the replay, it is recommended to restore the system after every replay when running multiple replays of the same captured workload.

The following two steps are necessary before replaying the captured workload:

1. Preprocess a captured workload
   The preprocessing step is required to optimize the captured workload file before replaying it. For more information, see *Preprocess a Captured Workload*.
2. Start the replayer
   The replay process is performed by the replayer, which should be running before starting the replay. The replayer is a service on operating system level that reads SQL commands from the preprocessed workload file and executes them one-by-one in timestamp-based order. For more information, see *Start and Stop the Replayer*.

You can preprocess a captured workload and replay the preprocessed workload using the *Replay Workload* card.

After replaying the preprocessed workload, you can generate replay reports. If you want to see the SQL statement parameters in the replay report, load the .cpt file after opening the workload analyzer.

Related Information

Preprocess a Captured Workload [page 547]
Start and Stop the Replayer [page 548]
Replay a Preprocessed Workload [page 550]
Replaying Configuration Settings [page 551]
Monitoring a Replayed Workload [page 554]
Generating Replay Reports [page 555]
Generate a Replay-Replay Comparison Report [page 556]
Set a Breakpoint [page 560]
Analyzing Replay Reports [page 557]
M_WORKLOAD_REPLAYS System View
M_WORKLOAD_REPLAY_PREPROCESSES System View

10.2.1.5.1 Preprocess a Captured Workload

The preprocessing step is necessary to optimize the captured workload file before replaying it.

Prerequisites

- You have the Control Replay Admin user for the control system.
- You have captured workloads using the *Capture Workload* card. For more information, see *Capture a Workload*.
- Copy the captured workload file from the source system to the target system in the trace directory. If you use a control system that is different from the target system, copy the captured workload file to the control system.
  You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

> → Recommendation
>
> Perform the pre-processing step in a separate system, not in the source system.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, choose the *Replay Workload* card.

   The *Replay Management* page opens displaying an overview of the captured workload located on the current system in the *Replay Candidate* tab.

2. Choose *Start* in the *Preprocess Status* column to start preprocessing the captured workload. If the status is not available, open the *Details* link to understand why.

Related Information

Capture a Workload [page 542]
System Privileges [page 538]

10.2.1.5.2 Start and Stop the Replayer

The replayer is a service on operating system level that should be running before starting the replay.

Prerequisites

- You have a user with the WORKLOAD REPLAY ADMIN system privilege to control the replayer. Store the logon credentials in the secure store. For more information, see *Secure User Store (hdbuserstore)* in the *SAP HANA Security Guide*.
- When using multiple replayers distribute and divide all involved components (for example, target instance, control instance, one or more replayers) on different hosts and systems.

> **i** Note
>
> The replayer is not a part of the SAP HANA database services that are running as deamon process. You must start and stop it yourself.

Procedure

1. Configure a hdbuserstore entry to authenticate the replayer with the database using the following command on the operating system on which your SAP HANA database is installed:

   ```
   hdbuserstore SET <key name> <host name@tenant database name> <user name> <password>
   ```

2. Start the replayer using the following command on the Linux command line of the system that you want to start the replayer on:

   ```
   hdbwlreplayer -controlhost <controlHost> (-controlport <controlPort> | -controlinstnum <controlInstanceNumber> [-controldbname <controlDatabaseName>]) -controladminkey <userName,secureStoreKey> -port <listenPortNumber>
   ```

> **Note**
>
> Running the command on the target system does not trigger the replay, it only starts the replayer.

The `controlhost`, `controlinstnum`, `controladminkey`, and `controldbname` parameters indicate the location of the control system.

Parameter Description

Parameter	Description
`controlhost`	Specifies the database host name of the control or target system (without a sqlport).
`controlinstnum`	Specifies the database instance number.
`controladminkey`	Specifies the user name and secure store key of the control management connection separated by a comma.
`controldbname`	Specifies the database name. When connected to a tenant, the tenant name should be used. When connected to a system database, the system database should be used as control database name.
`port`	Specifies the discretional port number for internal communication. You can use every port which is currently free. This port is used as long as the replayer is running. **Note** Do not to use any default ports or ports used by other processes (for example, 22 or 8080).
`controlport`	Specifies the control instance.

To start multiple instances of the replayer in parallel, define a specific port for each instance. The setup of the second instance fails when two instances run on the same port.

When running replays with a longer duration, you can add `&` at the end of the command line. This starts the process in the background and you can close the terminal connection immediately.

> **Example**
>
> ```
> hdbwlreplayer -controlhost <controlHost> (-controlport <controlPort> | -
> controlinstnum <controlInstanceNumber> [-controldbname
> <controlDatabaseName>]) -controladminkey <userName,secureStoreKey> -port
> <listenPortNumber> &
> ```

> **Note**
>
> If you want to use SSL encryption with the replayer, navigate to the wlreplayer.ini file on the OS and edit or add the section [replay_client] as follows:
>
> - Add the parameter `enable_target_ssl_connection = [true|false]` to enable SSL connections between the target system and the replayer

- Add the parameter `enable_control_ssl_connection = [true|false]` to enable SSL connections between the control system and the replayer

Only fully qualified domain names can be used for <controlHost> when starting the replayer.

3. If the console is still open, use Ctrl+C to stop the replayer. Alternatively, identify the OS process ID of the running replayer and shut it down using `kill<pid>`.

Related Information

Secure User Store (hdbuserstore)
SAP Note 2362820

10.2.1.5.3 Replay a Preprocessed Workload

You can replay all preprocessed workloads as often as necessary.

Prerequisites

- You have the privileges of a Target Replay Admin user. See *System Privileges* in the *SAP HANA Administration Guide for SAP HANA Platform*.
- The target system meets the same security and privacy prerequisites as the source system. Since the target system processes the same data as the source system, it should meet an appropriate security level depending on how critical your data is.

  → Recommendation

  Do not allow unnecessary network connections to the target system. Users registered on the source system could access the target system after the replay is completed.
- You have preprocessed the captured workloads using the *Replay Workload* card. For more information, see *Preprocess a Captured Workload*.
- The replayer is running. For more information, see *Start and Stop the Replayer*.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Replay Workload* card or click the *Start New Replay* link on the card.

The *Replay Management* page opens displaying the captured workload.

2. To change the preprocess destination, choose *Configure Replay* on the bottom right.

   After the preprocessing is completed, the preprocessed workload file is stored by default in the $SAP_RETRIEVAL_PATH/trace directory. Since the default trace directory generally resides in the same storage area with data and log volumes, preprocessing workloads may affect the performance across the entire system. If you enter a different destination, you may have a better distribution of the disk I/O between the data and log volumes, and the preprocessed files.

3. Choose a preprocessed workload with the status *Preprocessed* to start configuring it for the replay.

   The *Replay Configuration* page opens allowing you to configure various mandatory and optional settings. For more information about each setting, see *Replay Configuration Settings*.

   > **i Note**
   >
   > If a database backup is available, restore the database before starting the replay in the target system. For more information, see *SAP HANA Backup and Recovery*. When running a replay on a target system that has been restored using a backup taken automatically during the capture process, activate the *Synchronize Replay with Backup*.
   >
   > If no or only outdated database backups are available, you can still restore the database or manually export parts of the data before starting the replay in the target system. When running a replay on a target system that was restored using old backups or contains only smaller manual exports of data, deactivate the *Synchronize Replay with Backup* option.

4. After configuring the replay, choose *Review* to view the replay configuration.
5. To start the replay, choose *Confirm*.

   The *Replay Management* opens displaying the workloads that are being replayed in the *Replay List* tab.

Related Information

SAP HANA Database Backup and Recovery
Start and Stop the Replayer [page 548]
Replaying Configuration Settings [page 551]
Replay a Preprocessed Workload [page 550]
Analyzing Replay Reports [page 557]
System Privileges

10.2.1.5.3.1 Replaying Configuration Settings

The *Replay Configuration* allows you to configure several settings.

The *General Information* page allows you to customize the following mandatory and optional settings:

General Replay Information

Setting	Replay Name
Type	Mandatory
Description	Enter a name for the replay. By default this field has the same name as the initial captured file.

Setting	Description
Type	Optional
Description	Enter a description for the replay for your future reference. This information can be used when changing settings for different replays.

Target System Information

Setting	Host
Type	Mandatory
Description	Select a host name from the list. The information for the identifier, instance number, port number, and container will be filled automatically.

Setting	Identifier, Instance Number, Port Number
Type	Mandatory
Description	These fields are automatically filled when selecting the host.

Setting	Container
Type	Mandatory
Description	This field is automatically filled when selecting the host.

Target Instance Options

Setting	Synchronize Replay with Backup
Type	Optional

Setting	Synchronize Replay with Backup
Description	This option allows you to synchronize the replay with an existing database backup. It compares the log position ID of each transaction with the restart log ID of the backup used for recovery of the target system The option is turned on by default allowing the replayer to compare each statement with the database backup. This option makes it possible to check if there are no duplicate inserts and if the backup and replay are aligned. A backup is required for this option to work correctly. If it's tuned on without a recovery or with a wrong recovery, transactions might be rolled back. If the option is turned off, the replayer will replay statements even if no backup is present. If turned off and no recovery from the backup created during capture has been made, it's possible that transactions are missed during replay. If turned off, but recovery from the backup created during capture has been made, it's possible that unexpected results can occur during replay.

Setting	Collect Explain Plan
Type	Optional
Description	Collect the output of the EXPLAIN PLAN command for captured statements. You can use this information for comparison after the replay.

Setting	Transactional Replay
Type	Optional
Description	This option enables guaranteed transactional consistency during a replay. **i Note** Enabling this option may cause overhead to query runtime as transactional consistency needs to be checked constantly.

Optional Filter

Setting	Optional Filter
Type	Optional
Description	Use the *Optional Filter* to selectively replay only desired aspects of workloads. Filters can include different aspects such as *Application Name*, *Application User Name*, *Database User Name*, *Client*.

The *Replay Information* page allows you to customize the following mandatory settings:

Setting	Replayer Authentication
Type	Mandatory

Setting	Replayer Authentication
Description	Enter the database user who has the WORKLOAD REPLAY ADMIN privilege and will be used for the final preparation steps in the target instance. Select the authentication method and enter the credentials.

Setting	Replayer List
Type	Mandatory
Description	Check the running *Replayer* that will be used to connect to the target system and facilitate the replay.

Setting	User Authentication
Type	Mandatory
Description	Enter the SYSTEM user and the technical user. For a realistic replay, all users that are part of the workload, which has been chosen to be replayed, must be authenticated. For more information, see *Secure User Store (hdbuserstore)* in the *SAP HANA Security Guide*. To conduct a stable replay, reset the passwords of the database users. To reset the password for the database users captured in the source system you have the following options: • You can reset the passwords at once for all users not authenticated by external tools or created during the replay. As a result, these users will receive new passwords and get authenticated for the replay. Choose *Reset Password* to reset all passwords for all users except for the SYSTEM user and the technical user. This can be helpful when you don't know the actual password of each user. On the *Reset Password* window enable the confirmation box and choose *Confirm*. All selected user passwords in the defined target system will be changed as defined in this step. • You can authenticate users manually using the secure user store keys. The *User List* provides details on the authentication method for each user.

Related Information

Secure User Store (hdbuserstore)

10.2.1.5.3.2 Monitoring a Replayed Workload

You can monitor a replayed workload.

The *Replay Management* shows the workloads that are being replayed in the *Replay List* tab. You can start multiple replays in parallel. For more information on each replay, open the *Details* link in the message field.

To stop a replay when it is in progress, choose *Stop* in the *Replay Status* column. The status of the replay changes then to *Stopped*.

Import replay reports using the *Import Replay* button on the *Replay List* tab. For information on the security implications and configuration steps needed for importing replay reports, see *SAP Note 2109565*.

To access the *Replay Monitor*, choose the running replay. The monitoring view provides information such as duration, number of statements, size, and other details about the replay in progress. You can navigate away from the monitoring view using the arrow on the top right and can return anytime.

If you have already replayed preprocessed workloads, you can generate comparison reports for further analysis. For more information, see *Generating Replay Reports*.

You can monitor the preprocessed workload in the M_WORKLOAD_REPLAY_PREPROCESSES system view and the replayed workload in the M_WORKLOAD_REPLAYS system view. For more information, see *M_WORKLOAD_REPLAY_PREPROCESSES System View* and *M_WORKLOAD_REPLAYS System View* in the *SAP HANA SQL Reference Gudie*.

Related Information

Generating Replay Reports [page 555]
M_WORKLOAD_REPLAYS System View
M_WORKLOAD_REPLAY_PREPROCESSES System View
SAP Note 2109565

10.2.1.5.3.3 Generating Replay Reports

You can generate replay reports after successfully replaying a captured workload.

On the *Replay Management* page you can generate replay reports displaying:

- Capture-Replay Comparison
  You can open a capture-replay comparison choosing a replay from the *Replay List*. When opening a comparison report directly from the *Replay List*, the report shown always compares values from the original captured workload with values from the replay.
- Replay-Replay Comparison
  When using the *Compare Replays* button on the bottom right, the report shown always compares different replays with each other based on the same initial captured workload.

10.2.1.5.3.3.1 Generate a Replay-Replay Comparison Report

You can compare two or more replayed workloads with each other based on the same initial captured workload.

Prerequisites

- You have a user with the WORKLOAD REPLAY ADMIN, CATALOG READ, and WORKLOAD ANALYZE ADMIN system privileges.
- You have replayed preprocessed workloads using the *Replay Workload* card. For more information, see *Replay a Preprocessed Workload.*

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

You can start the comparison of the replayed workloads from the *Replay List* in the *Replay Management* page.

Procedure

1. On the *Database Overview*, with the *Administration* or *All* view selected, click the *Replay Workload* card.
2. Click the *Replay List* tab and then click *Compare Replays* on the bottom right.

   The *Select Baseline Replay* dialog opens allowing you to select the replayed workload that you want to compare. Use the *Target SID* information to distinguish between the replays.
3. Select one entry from the displayed list and click *Close*.

   The *Select Target Replay* dialog opens allowing you to select the replayed workload that you want to compare with the previously selected workload. The list displays replayed workloads based on the same initial captured workload.
4. Select one or more entries from the displayed list and click *Compare Replays* on the bottom right.

   The *Replay Report* opens displaying a comparison of the selected replayed workloads.

Related Information

System Privileges [page 538]
Replay a Preprocessed Workload [page 550]
Analyzing Replay Reports [page 557]

10.2.1.5.3.4 Analyzing Replay Reports

You can use comparison reports to analyze the completed replay.

The information is displayed on four tabs:

- Overview
- Load
- Performance Comparison
- Result Comparison

On the performance and result comparison tabs you can perform the following actions:

- Download detailed information in JSON format by choosing the download button at the top of the statement detail table.
  In the *Download SQL Statements* dialog, you can select the category of statements as well as the number of statements. For the performance comparison, these categories are *All*, *Comparable*, *Faster*, *Slower*, *Replay Failed*. For the result comparison, these categories are *All*, *Identical*, *Verification Skipped*, *Replay Failed*.
- Export replay reports to store them outside the database using the *Export Replay* button.

On the performance and result comparison tabs you can open the execution details for a specific statement by selecting the statement from the list. You can use the detailed execution level of both report types to compare the EXPLAIN PLAN results between the initial captured and replayed workloads, or between the baseline and target workloads. Comparing the plans can provide guidance for further statement-level investigation. This is only possible if the *Collect Explain Plan* setting was activated for both the capture and the replay during the configuration steps. For more information about this setting, see *Capture Configuration Settings* and *Replay Configuration Settings*.

Related Information

Replaying Configuration Settings [page 551]
Capturing Configuration Settings [page 543]

10.2.1.5.3.4.1 Overview

The *Overview* tab displays an overall comparison of the SQL statements involved in the capturing and replaying process.

> **i** Note
>
> When you turn off the *Transactional Replay* option on the *Replay Configuration* page, the result comparison does not include guaranteed transactional correctness. When running replays with this option turned on, the performance comparison can include a runtime overhead due to the transactional consistency.

Overview Information

Block Name	Description
Result Comparison	In a result-based comparison you get an overview of the statements with identical or different results. Click the block to open directly the *Result Comparison* tab.
Performance Comparison	In a performance-based comparison you get an overview of the statements based on a comparison of runtimes. You can change the tolerance ratio selecting a value from the drop-down list or entering a new value. Click the block to open directly the *Performance Comparison* tab.
Different Statements	Displays the top SQL statements that have different results from the selected baseline in descending order. You can click each row to open the *Execution Detail* page for the selected SQL statement. Use the drop-down arrow to filter the statements by time or by the number of records that have different results.
Slower Statements	Displays the top SQL statements that have a different performance ordered by the difference in execution time.Use the drop-down arrow to filter the statements by elapsed time, CPU time, or by execution time. You can click each row to open the *Execution Detail* page for the selected SQL statement. To view KPI details for each statement, you can click the icon on the right at the end of the row.
Verification Skipped	Displays the distribution of reasons for statements with skipped result comparison.
Replay Failed	Displays the distribution of reasons for the statements, which failed during replay. Use the drop down arrow to filter the statements by time or error code.
Capture Information	Displays information on the capture system, capture options, and the properties of the capture file.
Replay Information	Displays information on the replay system and the replay options. If the comparison was made between two replays, the information is displayed in a *Baseline Replay Information* block and in a *Target Replay Information* block.

Related Information

Replay a Preprocessed Workload [page 550]

10.2.1.5.3.4.2 Load

The load chart compares both the capture and the replay based on selected KPIs.

This tab includes load charts comparing both the captured and the replayed workloads after a capture-replay comparison, or the baseline and the target workloads after a replay-replay comparison. The capture values are represented by a solid line, the replay values are represented by a dotted line.

The KPIs can be toggled independently for both the capture and replay aspects making it easier to compare them with each other. Additional KPIs can be added using the *Show More KPIs* button on top right of the load chart.

10.2.1.5.3.4.3 Performance Comparison

The performance-based comparison provides an overview of statements compared by runtime.

Based on the selected tolerance ratio, the statements are classified as *Comparable* when they have a similar runtime within the defined tolerance ratio, *Faster*, *Slower*, or *Failed*.

You can sort the displayed statements by clicking directly the column header. You can sort them in ascending or descending order or you can filter the list. You can use the fields and buttons on top to:

- Search for a specific statement
- Change the tolerance ratio
- Display more statements
- Restore to default
- Refresh the displayed data
- Download detailed information in JSON format
- Change the sorting criteria
- Change filters
- Group
- Add more columns for analysis

To view a summary of how the execution time was spent, select the statistics symbol at the end of each statement line.

To open the execution details for a specific statement, select the statement from the list. On the *Execution Details* page, you can:

- Create a list of further statements, by choosing *Show More Statements* on the top right
- Search the list of individual executions
- View any parameters that might have been part of the query. To display the parameters, the associated captured workload file must be loaded in the workload analyzer.
- View the runtime *KPIs* for that execution
- Display the explain plan comparison for that statement

For failed executions you can see why the statement execution failed and what reasons can be investigated further by manual statement-level analysis.

10.2.1.5.3.4.4 Result Comparison

The result-based comparison provides an overview of statements with, for example, identical or different results.

The result-based replay report also includes a classification of statement types based on the content of those statements being either deterministic or non-deterministic. Deterministic statements should always deliver the

same results during a replay. Non-deterministic statements are expected to deliver different results (for example, because they don't contain an explicit sorting of results).

Statements are classified as:

- Identical if their result sets have the same row count and the same result hash
- Different if any of these criteria differ between capture and replay
- Skipped if they are related to system calls, monitoring view accesses or other internal actions, which don't require a result-based check
- Failed if they returned an error code

You can sort the displayed statements by clicking directly the column header. You can sort them in ascending or descending order or you can filter the list. You can use the fields and buttons on top to:

- Search for a specific statement
- Display more statements
- Restore to default
- Refresh the displayed data
- Download detailed information in JSON format
- Change the sorting criteria
- Change filters
- Group
- Add more columns for analysis

To open the execution details for a specific statement, select the statement from the list.

On the *Execution Details* page, you can search the list of individual executions, bring up a list of other statements to choose from with the *Show More Statements* button, view runtime KPIs for that execution as well as the result-based values, view any parameters that might have been part of the query, or display an explain plan comparison for that statement.

For skipped executions you can see why the verification of an execution was skipped.

For executions with different results you can see why the results are classified as being different. This classification is based on the result set's row count and result hash.

For failed executions you can see why the statement execution failed and what reasons can be investigated further by manual statement-level analysis such as the performance comparison tab.

10.2.1.5.3.5 Set a Breakpoint

When you identify a problem in the replay report, you can set a breakpoint to pause the replay exactly where the identified problem occurred.

Prerequisites

Context

You can start the replay a second time after setting breakpoints. When starting the replay again, this would allow you to look for example at system views before that statement is executed, or to look at the trace file.

You can set a breakpoint from the performance comparison or results comparison tabs on the *Replay Report* page.

Procedure

1. Open the *Execution Details* by choosing a statement from the SQL statements list.
2. Choose *Set Breakpoints* on the top right.
3. Choose *Save* on the top right.
4. Go back to the *Replay Management* page to configure and start the replay again.
5. After starting the replay, open the *Replay Monitor* to view the breakpoints and a list of all statements.

   On the bottom right you can use the *Resume Replay* button to execute the paused statement and continue to the next breakpoint. Or you can use the *Execute and Pause* button to execute the paused statement and then pause again.

10.2.2 Analyzing Workloads

Analyzing workloads from an SAP HANA database with the workload analyzer can help you identify the root cause of performance issues.

The following sections provide an overview of the workload analyzer tool:

What is the workload analyzer tool?

The workload analyzer is a tool that allows you to analyze the workload captured with the capture and replay tool, or any other workload occurring in a system.

The workload analyzer has two versions:

- The workload analyzer based on thread samples
  This version uses thread samples data to analyze the performance.
- The workload analyzer based on captured workloads
  This version of the tool performs the analysis using the captured workloads containing all the statistics of the query execution. Furthermore, it allows a deeper analysis with a full set of execution details.
  For more information on how to capture the workload of a system with the capture and replay tool, see *Capture a Workload*.

What is the workload analyzer based on captured workloads?

The workload analyzer based on captured workloads is a solution for analyzing database performances based on the workload captured with the capture and replay tool.

The workload analyzer based on engine instrumentation offers you two analysis types:

- On the upper part of the screen, the chart displays the system resource usage.
- On the lower part of the screen, three sections are displayed with further analysis:
  - *Capture Information*
    This section offers a detailed overview of the information on the captured workload, such as a summary, capture overview, filter options, database, and more. Selecting the section description will allow you to jump to the information you are interested in.
  - *SQL Statements*
    In this section, the graph displays the top SQL statements of the captured workload, and allows you to select the dimension you are interested in for graphical representation.
  - *Timeline*
    This section offers a timeline analysis based on an application and statement level hierarchy enables you to evaluate anomalies in the timeline and identify potentially problematic statements.
  - *Threads*
    In this section the main stacked area chart displays a more detailed visualization of the chart on the upper part of the screen based on a selected dimension (for example, thread type) over a given period of time. The bar charts located at the bottom next to it display the top five statements consuming most of the threads during the given timeframe. Clicking a specific statement hash opens a dialog with detailed statement information and an option to analyze it further with the SQL analyzer tool.

In contrast to the workload analyzer based on thread samples, using this tool requires to capture and load the workload before analyzing the performance. Capturing all the workload by default is not recommended because it introduces an overhead to system performance.

For more information on how to use the workload analyzer based on engine instrumentation, see *Analyze Workloads Based on Engine Instrumentation.*

What is the workload analyzer based on thread samples?

The workload analyzer based on thread samples is a solution for analyzing database performance using thread samples.

The workload analyzer based on thread samples provides a workload analysis using different KPIs, and it offers the following information sets:

- On the upper part of the screen, the chart displays the system resource usage. The chart displays both a real-time and a historical analysis. The information displayed on the grey background represents the historical analysis of the workload. Both analysis types are based on the sampling data. However, the historical analysis contains only aggregated data.
- On the lower part of the screen, the main analysis page offers the following four sections:
  - *SQL Statements*
    This section displays the analysis chart displaying the number of threads by lock wait time, and below, statement information and SQL statement. The option to navigate to the SQL analyzer tool for further investigation is provided by clicking on an entry on the chart.

- *Background Jobs*
  This section displays in the main chart information on the job progress. The miniature chart shows the system load data within the time range specified in the upper load chart, and the table below displays information on the delta merge.
- *Timeline*
  This section displays a timeline chart of the workloads, and a statements table containing thread-level information below. Clicking a block on the table will highlight the corresponding statement entry in the table.
- *Threads*
  In this section the main stacked area chart displays a more detailed visualization of the chart on the upper part of the screen based on a selected dimension (for example, thread type) over a given period of time. The bar charts located at the bottom next to it display the top five statements consuming most of the threads during the given timeframe. Clicking a specific statement hash opens a dialog with detailed statement information and an option to analyze it further with the SQL analyzer tool.

For more information on how to use the workload analyzer based on thread samples, see *Analyze Workloads Based on Thread Samples*.

Why should you use the workload analyzer tool?

The workload analyzer gives you an overview of the system's health at a glance. Moreover, the tool helps you identify the root cause of performance issues either by a real-time analysis or by reviewing historical data.

How to access the workload analyzer tool?

You can access both versions of the workload analyzer from the SAP HANA cockpit as follows:

- To access the workload analyzer based on thread samples, use one of the following:
  - From the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Analyze Workloads* on the *CPU Usage*, *Disk Usage*, or *Memory Usage* cards.
    or
  - From the *Database Overview* page, with the or *All* view selected, click *Monitor Performance* on the *CPU Usage*, *Disk Usage*, or *Memory Usage* cards. From the
    On the *Performance Monitor* page, click the ⌄ *(Navigation menu)* icon and choose *Workload Analysis*.
    or
  - From the *Database Overview* page, with the *Monitoring* or *All* view selected, click *Analyze Memory History* on the *Memory Usage* cards.
    On the *Memory Analysis* page, click the ⌄ *(Navigation menu)* icon and choose *Workload Analysis*.
- To access the workload analyzer based on captured workloads, use one of the following:

  > **i** Note
  >
  > Only workloads with the status *Loaded* can be analyzed. If the desired workload is *Not Loaded*, select *Start* next to the status to load the captured workload before proceeding.

  - On the *Workload Analysis* page, click *Capture*.

or

- From the *Database Overview* page, with the or *All* view selected, click the *Capture Workloads* card. The *Capture Management* page opens displaying a list of captured workloads.
- Click on a *Workload Analysis* status of the workload you wish to analyze. The *Capture Report* page opens where you can select to *Analyze Workload*.
- Open the workload analyzer based on thread samples. The chart on system resource usage on the top of the page displays areas highlighted in green when a workload has been captured. Drag across the highlighted area.

> **i Note**
>
> If a workload has not been loaded yet, you need to click on icon in the highlighted area and select *Load Capture*. The *Capture Report* page opens, allowing you to *Load Workload* on the top of the page.

Related Information

Capture a Workload [page 542]
Analyze Workloads Based on Thread Samples [page 564]
Analyze Workloads Based on Captured Workloads [page 566]
The Performance Monitor [page 515]
Memory Analysis [page 139]

10.2.2.1 Analyze Workloads Based on Thread Samples

You can analyze database performance using the workload analyzer based on thread samples.

Prerequisites

You have the system privileges `CATALOG READ` and `INIFILE ADMIN`.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. Open the *Workload Analysis* page. See *Analyzing Workloads*.

The workload analyzer opens displaying a chart on system database usage on the upper part of the screen and a more detailed visualization distributed in four sections:

View descriptions

View	Description
Top SQL Statements	A graphical representation of number of threads by lock wait time, and statement information and SQL statement.
Background Jobs	A chart displaying information on the jobs running in the background, and a delta merge table.
Timeline	A graph displaying the workloads by application name and statement hash, and a list of statements in a table.
Threads	Two graphs with customizable dimensions.

2. Analyze the displayed charts using the following features:
   a. On the top part of the screen, set the observed time range by selecting from presets or entering a custom time range, or by using the navigation buttons.
   b. In the settings menu, select desired KPIs. The selected KPIs appear in the legend area on the left-side of the chart. For a list of all available KPIs, see *Key Performance Indicators*.
   c. Customize your chart by applying a number of filtering options to only analyze the workloads you are interested in.
   d. Import and export datasets in order to store the data in an application and to analyze it in another system.
   e. Click the *(Navigate to other apps)* icon and choose *Performance Monitor* to monitor and analyze the performance of the SAP HANA database over time..
   f. Navigate to the *Current Table Distribution* page to see how tables are distributed across the hosts using one of the following:
      - In the *Top SQL Statements* view, left-click or Ctrl+left-click one or more statements from a chart or table. In the dialog window, choose *Open in Current Table Distribution*.
      - In the *Timeline* view, left-click or Ctrl+left-click one or more statements from the chart. In the dialog window, choose *Open in Current Table Distribution*.
      - In the *Timeline* view, in the Threads table, click the icon in the *Accessed Tables* column for a statement. In the dialog, choose *Open in Current Table Distribution*.
      - In the *Threads* view, left-click or Ctrl+left-click one or more bars from the secondary dimension chart. In the dialog, choose *Open in Current Table Distribution*.
   g. On the lower part of the screen in the *Top SQL Statements* view, select a SQL statement on the graph to see the statement hash, wait time, and number of waiting threads. Optional: To navigate to the *SQL Analyzer* page, click *Open in SQL Analyzer* to analyze query execution performance.
   h. On the lower part of the screen in the *Background Jobs* view, edit the granularity of the chart, and customize the delta merge table by setting the sorting order and defining the displayed columns.
   i. Customize the timeline chart in the *Timeline* view by adding new dimensions. Click *Edit* and select dimensions to add or remove from the chart.
   j. In the *Threads* view, see the workload analysis information you desire in the graphs by selecting from the primary and secondary dimensions on the left of each graph.

Related Information

Analyzing Workloads [page 561]
Key Performance Indicators [page 519]
Analyzing Statement Performance [page 568]
View or Modify a Table Distribution [page 489]

10.2.2.2 Analyze Workloads Based on Captured Workloads

You can analyze database performance using the workload analyzer based on captured workloads.

Prerequisites

You have the `WORKLOAD ANALYZE ADMIN system privilege`.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

In order to analyze the captured workload, the file has to be loaded into the database.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Capture Workload* card.

   The *Capture Management* page opens, displaying captured workloads that can be analyzed with the workload analyzer based on captured workload.

   > **i** Note
   >
   > Only workloads with the status *Loaded* can be analyzed. If the desired workload is *Not Loaded*, select *Start* next to the status to load the captured workload before proceeding. Please note that there is a 4 second delay in the initiation of the capturing process.

2. Click on the status of a loaded workload.

   The *Workload Analysis* page opens.

It displays a load chart overview of the workload on the top of the screen and the following information sections below:

Sections

Description
A summary of informative data on the captured workload.
A graphic representation of number of threads by lock wait time, and statement information and SQL statement.
A graph displaying the workloads by application name and statement hash, and a list of statements in a table.
Two graphs with customizable dimensions.

> **Note**
>
> This view is only available if the workload was captured in the same system.

3. Analyze the captured workload using the following features:
   a. Select the time range for the workload analysis by clicking on the date displayed on the top of the screen and selecting from a predefined set of time ranges or entering a custom range.
   b. Navigate back to the *Capture Management* page by clicking Captures.
   c. Customize the information displayed on the load graph on the upper part of the screen by selecting the host and services, filtering the results, or selecting the desired KPIs. The selected KPIs appear in the legend area on the left-side of the chart. For a list of all available KPIs, see *Key Performance Indicators*. Moreover, you can set specific filters (for example, statement hash, thread type, or application source) in order to analyze only the data you are interested in.
   d. On the lower part of the screen in the *Top SQL Statements* view, select a workload on the graph to see the statement hash, wait time, and number of waiting threads. Navigate to the *SQL Analyzer* page by clicking *Open in SQL Analyzer* to analyze query execution performance.
   e. Customize the timeline chart in the *Timeline* view by adding new dimensions. Click on *Edit* and select dimensions to add or remove from the chart.

Related Information

Capture a Workload [page 542]
Key Performance Indicators [page 519]
Analyzing Workloads [page 561]

10.2.3 Analyzing Statement Performance

Analyzing statement performance helps you understand performance issues of a query execution and other query execution aspects of the SAP HANA database.

The following sections provide an overview of the SQL analyzer tool:

What is the SQL analyzer tool?

The SQL analyzer is a query performance analysis tool of SAP HANA. The tool can be used to view detailed information on each query execution and can help you evaluate potential bottlenecks and optimizations for these queries.

How can you access the SQL analyzer?

You can open the SQL analyzer from the SAP HANA cockpit in one of the following ways:

- **Using the *SQL Statements* card**

  From the *Database Overview* page, with the *Monitoring* or *All* view selected, click*View all* on the *SQL Statements* card.

  Click the *Expensive Statements* tab.

  You can identify which SQL statements require a significant amount of time and resources. Each statement in the *Statement String*column has a *More* link, which opens the *Full SQL Statement* dialog.

  Click *Open in SQL Analyzer* to open the selected query with the SQL analyzer tool.

  Click *Open in SQL Analyzer* to open the selected query with the SQL analyzer tool.

  Click the *SQL Plan Cache* tab.

  Click the *More* link to open the *Full SQL Statement* dialog.

  Click *Open in SQL Analyzer*.

- **Using the *Manage SQL Performance* card**

  From the *Database Overview* page, with the *Administration* or *All* view selected, on the *Manage SQL Performance* card click *Statement Hints*, *Plan Stability*, *Plan Trace*, or *Saved Plans*.

  On each page, each statement string has a *More* link in the *Statement String* column, which opens the *Full SQL Statement* dialog.

  Click *Open in SQL Analyzer* to open the selected query with the SQL analyzer tool.

Using the *Workload Analysis* (based on thread sampling)

See *Analyzing Workloads Based on Thread Samples*.

- Using *Database Explorer*

  On the *Database Overview* page, click *Open SQL Console*. The SAP HANA database explorer opens.

  Click the *Analyze* button, and then click *Analyze SQL* to open the SQL analyzer.

- From the SAP Web IDE for SAP HANA you can open the SQL analyzer in a SQL Console by clicking *Analyze SQL* in the menu of the *Analyze* button.

Which views are supported by the SQL analyzer?

The following views are supported by the SQL analyzer:

SQL Analyzer Views

View	Description
Overview	This view provides an overview of the query execution including metadata for the analysis. It displays the following KPIs: • Time • Dominant Operators • Statistics • SQL Performance Recommendations To get more information on the KPIs, see *Key Performance Indicators*.
Plan Graph	This view provides graphical guidance to help you understand and analyze the execution plan of a SQL statement. In case of SQLScript, the SQLScript definition is also displayed.
SQL	This view displays the complete SQL string that is analyzed.
Operators	This view provides a list of operators used during query execution and includes additional details about each operator, which can be used for analysis.
Timeline	This view provides a complete overview of the execution plan based on the visualization of sequential time-stamps.
Tables in Use	This view provides a list of tables used during query execution and includes further details on tables, which can be used for further analysis. It can be used to understand which tables are needed to fulfill a given SQL statement execution.

View	Description
Table Accesses	This view provides details on the table accesses performed during the processing of a SQL statement, which can be used for analysis.
Compilation Summary	This view provides details on the query compilation process. It can be used to understand how much time was spent on which operation, which cost-based information was used, and what the plan properties are.
Recommendations	This view shows details on the provided recommendation allowing you to understand the reasoning for it.

Related Information

Statement Performance [page 570]
Monitor and Analyze Expensive Statements [page 529]
Monitor and Analyze Statements with SQL Plan Cache [page 528]
Statement Hints [page 590]
SQL Plan Stability [page 588]
Manage Saved Plans [page 578]
Analyze Workloads Based on Thread Samples [page 564]
Monitor and Analyze Statements with Plan Trace [page 573]

10.2.3.1 Statement Performance

The SQL analyzer is used to analyze statement execution performance.

Context

The *SQL Analyzer* can be opened from the SAP HANA cockpit or the SAP HANA database explorer. You can also open it from the SAP Web IDE for Full-Stack Development. For more information on how to open the SQL analyzer, see *How can you access the SQL analyzer?* in *Analyzing Statement Performance*.

Procedure

1. Open the SQL analyzer.

The *SQL Analyzer* page opens, displaying the following views:

- Overview
- Plan Graph
- SQL
- Operators
- Tables in Use
- Table Accesses
- Compilation Summary
- Recommendations

2. Open the *Overview* tab to see important KPIs required to begin a performance analysis before going into the complex details.

KPI Description

KPI	Description
Statistics	System version: The version of the system where the execution occurred Tables in use: Total number of tables touched by any operators during execution You can navigate to the *Tables in Use* tab on the table below by clicking this line. Result Records: The final result record count Memory Allocated: Total memory allocated for executing the statement Distribution: Number of SAP HANA index servers that are related to the query execution The name of the host
SQL Performance Recommendations	Recommendations on how you can improve the performance of SQL-related operations.

3. Open the *Plan Graph* tab to understand and analyze the execution plan of an SQL statement. It displays a visualization of a critical path based on inclusive execution time of operators and allows you to identify the most expensive path in a query execution plan.

   In case of a SQLScript, the *Plan Graph* displays its complete definition. To retrieve the information in a text format, you can copy the definition by clicking the copy icon.

   In the *Plan Graph* tab you can open the *Detail Properties* view by clicking one of the operators. This view offers detailed information on the operator such as name, location, ID, summary.

   You can open the edge information detail by clicking one of the links between the operators. This view offers further information on the edge values, such as target, source, output cardinality, fetch call count, and estimated output cardinality.

   Furthermore, you can configure plan graph settings. You can set the color of the nodes by type or location, and choose to show either physical or logical inner plans.

4. Open the *SQL* tab to get the complete view of the SQL statement string that is being analyzed.

5. Open the *Operators* tab to pinpoint specific operators of interest.

   The view lists characteristics of all operators and supports:
   - Display of various KPIs, for example physical (whether an operation is a real, physically executed one), offset, execution time, CPU time
   - Setting of filters along all the columns/KPIs
   - Display of the number of operators within the filtered set
   - Immediate aggregated information (max, min, sum, and so on) for the same KPIs on the filtered operator set
   - Detailed display of all operators within the filtered set
6. Open the *Tables in Use* tab for an overview of which tables have been used during the processing of a statement.

   The view displays the following information:
   - The name of the table
   - The host and port of the table's partition in the *Location* column
   - The partition number of the table's partition in the *Partition* column
   - The maximum possible amount of data processed for a given table during the execution of a statement, including the possibility of multiple accesses, in the *Max. Entries Processed* column
   - How often a table has been accessed during statement execution in the *Number of accesses* column
   - The maximum processing time across the possibly multiple table accesses in the *Maximum processing time* column
7. Open the *Table Accesses* tab to see the details on the table accesses performed during the processing of a statement.

   The view displays the following information:
   - The *Offset* time for accessing the table
   - The name of the table
   - The *Conditions* that affect the table accesses
   - The *Processing Time*
   - The amount of entries processed during an operation in the *Entries Processed* column
   - The host and port of the table's partition in the *Location* column
   - The operator's unique ID in a query execution result in the *Operator Id* column
   - The name of the operator
   - The detail information of an operator in the *Details* column
   - The partition number of the table's partition in the *Partition* column

   For cached plans, the table displays the following information: *Inclusive Estimated Cost*, *Exclusive Estimated Cost*, and the *Estimated Output Size*.

   Optional: To get aggregated information for each column, choose the aggregator functions in the drop-down menu under the column name. If an aggregator is chosen, you can see *More* information for your chosen query.

   Optional: To refine results, click on *Filters* on the upper right corner. The *Filters* page appears, where you can choose the information type, operator, and enter values according to which you wish to filter the *Table Accesses*. Choose *OK* to confirm. If you want to remove the filters you selected, go to the *Filters* page and choose to *Restore* to defaults.

   Optional: To sort the results, click the sorting icon next to the *Filters* button and choose the sorting order.

Optional: To customize the columns displayed in the information list, click on the settings icon on the upper right corner and choose to hide or display the desired columns.

8. Open the *Compilation Summary* tab to view details on the query compilation process, including compilation time breakdown, cost-based optimization details, and plan properties.
9. Open the *Recommendations* tab to see the details for each provided recommendation. This allows you to easier understand the reasoning behind the recommendation for SQL query optimization.
10. Optional: Re-execute the SQL query by clicking the *Re-execute* button on the top right corner. If the query is parameterized, you can change the parameter values.

    With a parameterized query, the *Input Parameters* appears, prompting you to enter your desired parameters in the *Parameter* tab. You can list the parameter values and separate them by a comma, or enter and edit them individually. Where applicable, you can check the *Empty value* box. The values you entered are reflected in the SQL string that you can see in the *SQL Statement* tab. Click on *Execute SQL* to analyze the SQL string according to the parameter values you entered and see the result on the *SQL Analyzer* page.

Related Information

Analyzing Statement Performance [page 568]

10.2.3.2 Monitor and Analyze Statements with Plan Trace

Plan Trace is a trace feature for SQL analyzer.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Plan Trace enables you to collect SQL queries and their execution plans, executed in a given time frame. For each SQL query traced you can visualize the execution plan for performance analysis. Only 'SELECT' statements are traced with Plan Trace.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *Plan Trace* on the *Manage SQL Performance* card.

   The *Plan Trace* page opens, displaying a set of statistics for each SQL statement collected in a given time frame. To find the most expensive SQL statements use the displayed categories (for example, start and end time, schema, user, statement hash).
2. Click the *Configure Trace* button on the bottom right to configure the plan trace.

   The *Configuration* dialog opens, allowing you to set options like maximum disk or memory usage.
3. Optional: Open the selected statement string with the SQL analyzer by clicking the *More* link in the *Full SQL Statement* dialog.

10.2.3.3 Analyzing SQL and Saving Plans

This function of the SAP HANA cockpit allows you to save SQL plans, download them, and load a file from the SAP HANA database.

There are two ways in which you can record SQL plans for future reference with the SAP HANA cockpit:

- You can **save** them within the SAP HANA database filesystem, and
- **Download** the saved SQL plans as local PLV files on your personal computer.

Both of these options are available with this functionality.

Additionally, you can also **load** PLV files into the SQL Analyzer tool from the trace folder in the SAP HANA database.

Related Information

Save Plans as Files [page 575]
Load Files [page 576]
Download Files [page 577]

10.2.3.3.1 Save Plans as Files

You can save SQL plans within the SAP HANA database.

Prerequisites

- Since this functionality can be accessed from different performance management and monitoring tools within the SAP HANA cockpit, please check the privileges they might require respectively.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Saving SQL plans helps you record them as PLV files in the trace folder within the SAP HANA database.

You can save SQL plans from the following performance monitoring pages:

- Expensive Statements
- Plan Stability
- Statement Hints
- SQL Analyzer

Procedure

1. On the *Database Overview* page, select the desired page.

Database Overview View	Card	Select
Administration or All	*Manage SQL Performance*	○ *Statement Hints* ○ *Plan Stability*
Monitoring or All	*SQL Statements*	○ *SQL Plan Cache* ○ *Expensive Statements*

   a. Click *More* beside the desired statement string.
   b. Choose *Analyze SQL and Save Plan*.

   or

   On the *SQL Analyzer* page click *Execute and Save* option from the *Execute* menu. Choosing *Execute* only will not save the plan.

2. Enter a prefix for a file name.
3. Choose whether to run the analysis as a background activity.

Related Information

Analyzing SQL and Saving Plans [page 574]
Monitoring and Analyzing with the Statements Monitor
Managing Plan Stability
Analyzing Statement Performance [page 568]
Managing Statement Hints
Statement Hints [page 590]
Download Files [page 577]

10.2.3.3.2 Load Files

You can load saved files within the SAP HANA database.

Prerequisites

- To load files, you have to have saved SQL plans in the SAP HANA database. For more information, see, *Save Plans*.
- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Loading saved SQL plans into the database will allow you to open the saved SQL plan with the *SQL Analyzer*.

You load saved SQL plans from the *Saved Plans* page.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Saved Plans* section on the *Manage SQL Performance* card.
2. Click *Load PLV*.

   A dialogue appears, prompting you to select files you want to load into the current system. The location of the files appears on the top of the file list.
3. Choose the desired files that you want to load into the SAP HANA database and confirm.

Related Information

Analyzing SQL and Saving Plans [page 574]
Save Plans as Files [page 575]
Analyzing Statement Performance [page 568]

10.2.3.3.3 Download Files

You can download saved SQL plans as local files to your hard drive.

Prerequisites

- You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.
- To download files, you have to have saved SQL plans in the SAP HANA database. For more information, see, *Save Plans*.
- Since this functionality can be accessed from different performance management and monitoring tools within the SAP HANA cockpit, please check the privileges they might require respectively.

Context

Downloading SQL plans helps you share the SQL plans as a PLV file with others.

You can download SQL plans on the *Saved Plans* page and the *SQL Analyzer* page.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click *Saved Plans* on the *Manage SQL Performance* card.
2. Download SQL plans as local files to your hard drive from the *Saved Plans* tool by choosing the download icon next to the desired statement string.
3. Download SQL plans as local files to your hard drive from the *SQL Analyzer* tool by clicking the *Download* button on the upper right corner of the executed parameterized query.

   A dialogue appears that allows you to *Download* the file to your hard drive.

Related Information

Analyzing SQL and Saving Plans [page 574]
Save Plans as Files [page 575]
Manage Saved Plans
Analyzing Statement Performance [page 568]

10.2.3.4 Manage Saved Plans

The SQL analyzer result page shows SQL plans saved from a previously executed query.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

The saved plans feature allows you to revisit SQL statement queries that were executed with the SQL analyzer in a previous session without having to re-execute them in the current session.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click on the *Saved Plans* section on the *Manage SQL Performance* card.

   The *SQL Analyzer result* page opens, displaying the *Saved Plans* table, containing the collected information on previously executed SQL queries. To find the your desired SQL statements, use the displayed categories (for example, system version, statement string, plan type, user name, schema name, statement hash, and so on).
2. Optional: You can delete saved plans in two ways:
   a. Mark the checkbox on the left of the listed statement, and select *Delete* on the top of the table.

      The *Delete Plan* dialog opens, asking you if you want to delete both the plan and the corresponding PLV file, or just the plan.
   b. Select a statement and click on the trash bin icon next to the statement string.

      The *Delete File* dialog opens, making sure if you want to delete the chosen PLV file.

> **i Note**
>
> This option does not get rid of the plan, just the PLV file. If you want to delete the plan as well, choose the former method.

3. Optional: You can search saved plans by entering a statement string keyword in the search field.
4. Optional: You can sort the order of the table and filter results by the desired parameter.
5. Optional: You can customize what columns are shown in the table in the settings menu.

Related Information

Analyzing Statement Performance [page 568]
Statement Performance [page 570]

10.2.4 Analyzing Memory with the Memory Profiler

Analyzing the memory consumption of the SAP HANA database over time can help you pinpoint bottlenecks, identify patterns, and forecast requirements. Use the *Memory Profiler* to record and visually analyze memory consumption data.

What is the memory profiler?

The memory profiler allows you to record the memory allocations and deallocations in SAP HANA services, such as the `indexserver` or the `nameserver`. The recording is carried out within a limited time frame which may range from a few minutes to several hours. Once a recording is stopped, the allocator activities can be analyzed based on the recorded data. The analysis of the collected allocator data reveals detailed insights into the memory situation of the SAP HANA system.

Use the memory profiler to identify and analyze the memory consumption in a system. It can provide answers to questions, such as:

- Which allocators consumed the most or the least memory?
- Which allocators performed the most or the least memory allocations/deallocations?
- Which allocators caused the highest peaks in memory consumption?
- Which allocators are not in use since they performed no allocations at all?
- Which allocators are responsible for memory leaks as they allocate more memory than deallocate?

How does the memory profiler work?

The memory profiler can only analyze recorded allocator data. You can create a recording for your system or import a recording that was created on another system.

The collected data is sampled data, not continuous data. The services of the SAP HANA system fetch the current data of the specified allocators periodically. By default, the sampling interval is 20 ms. You can change this value when starting a new recording.

> **i Note**
>
> The smaller the sampling interval, the more data is collected. Since the data is compressed, the increase in size is not linear.

As long as the recording is running, the collected data is kept in internal data structures in the main memory of the SAP HANA service. Once the recording stops, the data is written into a single trace file.

Trace files are stored in the trace directory of the SAP HANA tenant database or system database. The filenames follow the default naming convention:

```
<service>_<host>.<port>.memory.<recording name>.trc
```

> **i Note**
>
> Although the trace file extension is `.trc` and the files are stored in the trace directory, the memory profiler does not create regular SAP HANA trace files. Trace file rotation does not affect trace files created by the memory profiler.

The data in the memory profiler trace file must be uploaded to an SAP HANA database in order to by analyzed. Once you have uploaded the file to the database, it will undergo multiple compression runs. You will therefore notice changes in size as the data is optimized.

Related Information

Record Memory Allocation Data [page 580]
Analyze Memory Allocation Data [page 582]
Traces

10.2.4.1 Record Memory Allocation Data

Use the *Memory Profiler* page to record the allocation and deallocation of memory in the database. A recording can be visually analyzed to identify performance issues.

Prerequisites

- You have a user in the system database with the system privilege `MEMORY PROFILER ADMIN`.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview*, with the *Monitoring* or *All* view selected, click the *Profile Memory* link on the *Monitoring* card to access the *Memory Profiler*.
2. Choose *New Recording*.
3. Enter the required information on the *New Recording* page and select the allocators you want to record. You can choose for every allocator if its call stack is included in the recording.

Option	Description
Name	The name of the recording.
Host and Service	The host and service that you want to record. The selectable services depend on the host type.
Sampling Interval	The sampling interval of the recording. The shorter the interval, the more data is collected. The default value is 20 ms.
Stop Recording after	The time after which a recording stops automatically. The default value is 300 s.
Replace existing trace file	The recording overwrites any existing recording with the same name.
Automatically upload recording to database	The recorded data is uploaded to the database automatically after the recording has finished.

4. Choose *Start Recording* to start the recording process.

> **i** Note
>
> Only one recording can be created at a time.

Next Steps

You can monitor the recording status in the *Running Recordings* section. You can stop or cancel the recording for each of the recorded services individually. Choose *Stop* to stop the recording but keep the trace file. Choose *Cancel* to abort the recording and delete any recorded data.

Once the recording has finished, it is listed in the *Available Recordings* section and can be analyzed. For more information, see *Analyze Memory Allocation Data*.

Related Information

Analyze Memory Allocation Data [page 582]
Delete a Memory Profiler Recording [page 583]

10.2.4.2 Analyze Memory Allocation Data

Use the *Memory Profiler* page to analyze the allocation and deallocation of memory in the database.

Prerequisites

- You have a user in the system database with the system privilege `MEMORY PROFILER ADMIN`.
- You have created or uploaded a memory profiler recording.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview*, with the *Monitoring* or *All* view selected, click the *Profile Memory* link on the *Monitoring* card to access the *Memory Profiler*.
2. To analyze a recording, choose *Analyze* for the recording you want to analyze in the *Available Recordings* section.

   If the recording trace file you want to analyze is not available under *Available Recordings*, choose *Upload* to add it. The trace file must be located inside the trace file directory:

   ```
   /usr/sap/<SID>/HDB<instance>/<host>/trace/<db_name>
   ```

   Uploading a recording to the database may take some time. You can monitor the upload status in the *Running Uploads* section.

Results

The memory profiler opens the *Memory Recording Analysis* page. You can analyze the memory allocation and deallocation in a number of views.

View	Description
Summary	Use the *Summary* view to identify allocators with memory peaks, a high number of allocations, memory deltas, and the most allocated memory.
Timeline	The *Timeline* view provides a graphical representation of the recorded memory consumption. Select up to 10 allocators from the *Allocators* table to analyze the recorded memory allocations. You can select data points of a graph to analyze specific changes of the allocated memory. If a call stack was recorded for the allocator, select an allocation or deallocation to view the call stack.

View	Description
Bottom-up	The *Bottom-up* view allows you to drill down through a call stack to identify the source of memory allocation in the software.
Sizes	The *Sizes* view contains a chart depicting memory allocations, deallocations, and remaining allocations grouped by size. You can narrow down the data by selecting fewer top allocators.

Next Steps

You can switch to a different recording by selecting it from the *Recording* dropdown.

Related Information

Record Memory Allocation Data [page 580]
Delete a Memory Profiler Recording [page 583]

10.2.4.3 Delete a Memory Profiler Recording

You can use the cockpit to delete a memory profiler recording.

Prerequisites

- You have a user in the system database with the system privilege `MEMORY PROFILER ADMIN`.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview*, with the *Monitoring* or *All* view selected, click the *Profile Memory* link on the *Monitoring* card to access the *Memory Profiler*.
2. To delete a recording from the database, select it from the *Available Recordings* list. Then select the delete icon.

   The recording is deleted from the database. However, the recording trace file is still located in the trace file directory and can still be uploaded to the database to be analyzed.

3. Choose *Upload* to open the *Upload Recorded Data* page.
4. To delete a recording from the trace file directory, select it from the *Available Data in Trace Files* list. Then select the delete icon.

Results

The trace file is deleted.

Related Information

Record Memory Allocation Data [page 580]
Analyze Memory Allocation Data [page 582]

10.3 Improving Performance in SAP HANA Cockpit

You can improve the performance of the database using the SAP HANA cockpit.

You can use the following tools to analyze fine-grained aspects of system performance in the SAP HANA cockpit:

- Use the *Recommendations* card to get suggestions on how to improve and optimize your database.
- Use the *Data Cache* to monitor and manage different types of cached queries.
- *Manage plan stability* to restore performance speed from the previous to the current system.
- *Manage statement hints* to add statement hints to an SQL statement without modifying the actual statement in the application.

Related Information

Recommendations [page 470]
Data Cache [page 585]
SQL Plan Stability [page 588]
Statement Hints [page 590]

10.3.1 Data Cache

Use Data Cache for an overview and management options of different types of cached queries.

Cached data helps improve the performance of the SAP HANA database by allowing you to retrieve data quickly without repeated query execution.

On the SAP HANA cockpit overview page, the *Manage SQL Performance* card contains a link to the data cache page and shows information on the count of cached files.

The *Data Cache* page offers three views based on three different types of cached data:

- *Procedure Result Cache* allows you to monitor and manage cached intermediate results of table variables within SQLScript.
- *Static Result Cache* allows you to monitor and manage cached SQL result views and calculation views.
- *Dynamic Result Cache* allows you to monitor and manage up-to-date SQL query results of cached views.

Related Information

Monitor and Manage Procedure Result Cache [page 585]
Monitor and Manage Static Result Cache [page 586]
Monitor and Manage Dynamic Result Cache [page 587]

10.3.1.1 Monitor and Manage Procedure Result Cache

Monitor and manage procedure result cache.

Prerequisites

You need the following privileges for specific functions of procedure result cache:

- For system and monitoring views, you need the SELECT privilege (the PUBLIC role has this privilege by default for public and related views)
- For the CREATE/ALTER PROCEDURE, you need the OPTIMIZER ADMIN privilege.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Procedure Result cache is a view available on the data cache page.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Data Cache* section on the *Manage SQL Performance* card.

   The Data Cache page opens, where procedure result cache is the default view. You can see a table of the cached entries or switch to a view of variables the cached entries are related to.
2. Optional: Enable or disable data cache collection in the *Variables* tab.
3. Optional: Remove a cached entry or drop the variable with all related entries.
4. Support monitoring and management of procedure result cache in the following ways:
   - Search the results by object name or schema
   - Filter and sort the results by choosing the header of a column
   - Select which columns you want to be displayed and their order in the settings menu
   - Customize column width by adjusting the borders.

Related Information

Data Cache [page 585]
Monitor and Manage Static Result Cache [page 586]
Monitor and Manage Dynamic Result Cache [page 587]

10.3.1.2 Monitor and Manage Static Result Cache

Monitor and manage static result cache.

Prerequisites

You need the following privileges for specific functions of static result cache:

- For system and monitoring views, you need the SELECT privilege (the PUBLIC role has this privilege by default for public and related views)
- For the CREATE/ALTER PROCEDURE, you need the OPTIMIZER ADMIN privilege.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Static result cache is a view available on the data cache page.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Data Cache* section on the *Manage SQL Performance* card.

   The Data Cache page opens, where variable cache is the default view.
2. Choose the *Static Result Cache* view.

   The static result view opens, displaying a table of cached SQL query result entries. You can choose to see metadata the cached entries are related to, or to exclusions from static result cache.
3. Optional: Remove a cached entry or drop the metadata with all related entries.
4. Support monitoring and management of static result cache results in the following ways:
   - Search the results by object name or schema
   - Filter and sort the results by choosing the header of a column
   - Select which columns you want to be displayed and their order in the settings menu
   - Customize column width by adjusting the borders.

Related Information

Data Cache [page 585]
Monitor and Manage Procedure Result Cache [page 585]
Monitor and Manage Dynamic Result Cache [page 587]

10.3.1.3 Monitor and Manage Dynamic Result Cache

Monitor and manage dynamic result cache.

Prerequisites

You need the following privileges for specific functions of dynamic result cache:

- For system and monitoring views, you need the SELECT privilege (the PUBLIC role has this privilege by default for public and related views)
- For the CREATE/ALTER PROCEDURE, you need the OPTIMIZER ADMIN privilege.

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

Dynamic result cache is a view available on the data cache page.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click the *Data Cache* section on the *Manage SQL Performance* card.

   The Data Cache page opens, where variable cache is the default view.
2. Choose the *Dynamic Result Cache* view.

   The dynamic result cache view opens, displaying a table of the cached SQL query result entries. You can choose to see metadata the cached entries are related to, or to exclusions from dynamic result cache.
3. Optional: Remove a cached entry or drop the metadata with all related entries
4. Support monitoring and management of dynamic result cache results in the following ways:
   - Search the results by object name or schema
   - Filter and sort the results by choosing the header of a column
   - Select which columns you want to be displayed and their order in the settings menu
   - Customize column width by adjusting the borders.

Related Information

Data Cache [page 585]
Monitor and Manage Procedure Result Cache [page 585]
Monitor and Manage Static Result Cache [page 586]

10.3.2 SQL Plan Stability

SQL plan stability helps ensure the fast performance of queries by capturing query plans in a source system and reusing them in a target system to regenerate the original query plan.

In SAP HANA, the SQL query processor parses SQL statements and generates SQL query execution plans. As the query processor and the query optimizer continue to be developed (in, for example, the SAP HANA Execution Engine - HEX) the resultant execution plans for a given query may change from one SAP HANA revision to another. Although all developments are intended to improve performance of a query, it is possible it might not be equivalent after an upgrade.

In order to guarantee the performance of a query in new system upgrades, the plan stability feature offers the option to preserve a query's execution plan by capturing an abstraction of the plan and reusing it after the upgrade to regenerate the original plan and retain the original performance. In some cases, using statement hints may provide a solution to a loss of performance, see *Managing Statement Hints*.

Restrictions related to Plan Stability are documented in SAP Note 2639193: *SAP HANA SQL Plan Stability*.

Related Information

Manage SQL Plan Stability [page 589]
Statement Hints [page 590]
SQL Plan Stability
SAP Note 2639193

10.3.2.1 Manage SQL Plan Stability

Use the *SQL Plan Stability* app to capture query plans in a source system and reuse them in a target system to regenerate the original query plan.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview* page, with the *Administration* or *All* view selected, click on the *SQL Plan Stability* section of the *Manage SQL Performance* card.

   The *SQL Plan Stability* page opens. If you have any captured abstract SQL plans from the source system, they are displayed in the table view.
2. To capture SQL query plans from the source system, turn the *Capture Abstract SQL Plans* on.

   You are prompted to configure the capture options. Choose from registered filters or create your own, choose whether you want to include plans form the SQL Plan Cache, and select *Start Capture*.

   The *Capture Status* appears on the top of the screen, displaying the progress of captured execution plans and plans from SQL Plan Cache.

   Optional: Press the refresh button next to the *Capture Status* to see the captured plans in the table below.

   After the process is turned off, the captured query plans are displayed in the *Captured Plans* table.

   Optional: If you want to terminate the capture process prematurely, turn the *Capture Abstract SQL Plans* off.
3. Optional: See the full SQL statement string of a query by selecting *More* next to the visible statement string snippet in the table.

4. To apply the imported SQL query plans to the current system, turn the *Apply Abstract SQL Plans* on.

   > **i Note**
   >
   > This will only apply to the SQL queries listed in the *Captured Plans* table.

   Optional: Enable or disable the SQL query plans being applied to the current system by selecting them in the table and choosing *Enable* or *Disable* above the table. You can also *Select All ASP* query plans, or choose to *Delete* plans.

   Optional: To see the progress of the query plan application, press the refresh button.
5. Optional: You can adjust the settings such as the maximum number of saved plans and maximum memory allocation in the *Configure Capture* dialog.
6. Optional: You can change the location details (host and port) of the server by selecting *Update Locations* from the *Advanced* options menu.

   The *Update Locations* dialog opens where you can set the desired host and port.
7. A migration of abstract SQL plans is necessary after a system upgrade. If you are trying to use Plan Stability features with outdated data, you will be prompted to migrate the abstract SQL plans into the upgraded system. You can also do this actively by by selecting *Migrate ASPs* from the *Advanced* options menu.

   The *Migrate Abstract SQL Plans* dialog opens, prompting you to confirm your choice.

   > **i Note**
   >
   > During the migration of abstract SQL plans to the upgraded system other Plan Stability functionalities are not available.

10.3.3 Statement Hints

Use *Statement Hints* to add statement hints to an SQL statement without modifying the actual statement in the application.

Statement Hints allow you to pair an SQL statement string with a string of hints to be used during execution. Whenever a particular SQL statement is then executed in SAP HANA, the assigned statement hints are automatically added to the statement for execution.

Open the *Statement Hints* page by clicking on the *Statement Hints* section of the *Manage SQL Performance* card.

11 Backup and Recovery

SAP HANA offers comprehensive functionality to safeguard your database and ensure that it can be recovered speedily and with maximum business continuity. Use SAP HANA cockpit to manage backups, and recover a database.

SAP HANA cockpit supports the following backup and recovery capabilities:

- Manual and automated backups
- Extensive configuration options
- Backup lifecycle management (housekeeping)
- Recovery to a point-in-time
- Recovery to a specific full backup or data snapshot
- Database copy using backup and recovery
- Full integration with third-party backup tools

Administering Databases in SAP HANA Cockpit

To administer a database, the database must first be registered as a database in SAP HANA cockpit.

For more information, see *Set up SAP HANA Cockpit for the First Time* in *SAP HANA Administration with SAP HANA Cockpit*.

Related Information

Set Up SAP HANA Cockpit for the First Time [page 10]
Display the Backup Configuration Settings [page 592]
Change the Backup Configuration Settings [page 592]
Create Data Backups [page 603]
Create a Data Snapshot (Native SQL) [page 606]
Schedule Backups [page 611]
Recover a Database [page 628]
Copy a Database [page 640]
Housekeeping: Deleting and Archiving Backups [page 621]

11.1 Display the Backup Configuration Settings

Using SAP HANA cockpit, you can display an overview of the active backup configuration settings for a specific database.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview*, with the *Administration* or *All* view selected, click the *Database Backups* card.
2. One the card, click *Backup Configuration*.

   The configuration settings for the database are displayed.

   > **i** Note
   >
   > On this page, most configuration settings are only displayed. For the **current database**, you can **edit** the configuration settings for the Backup Scheduler and the backup retention policy.

Related Information

Change the Backup Configuration Settings [page 592]

11.2 Change the Backup Configuration Settings

Using SAP HANA cockpit, you can change the default backup configuration settings for the system database and all the tenant databases in the SAP HANA system.

Prerequisites

You have the authorization DATABASE ADMIN.

Context

To change the default configuration settings for all the databases in an SAP HANA system, you need to work through the system database.

Procedure

1. From the *Database Directory*, select a system database and choose *Database Management*.

   An overview of the databases in the SAP HANA system is displayed.
2. Choose *Backup Configuration*.

   An overview of the current systemwide default backup configuration settings is displayed.

   The backup encryption status for the SAP HANA system is displayed. For more information, see *SAP HANA Backup Encryption* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.
3. To change a group of configuration settings, navigate to that group and choose *Edit*.

   When you edit a settings group, SAP HANA cockpit checks whether the configuration settings are within the recommended range. If a setting has been changed to a non-recommended setting, a setting within the recommended range is proposed. You then have the option to save or discard the recommended setting.
4. Save or discard your changes.

   To reset a group of configuration settings to the default values, choose *Reset to Default*.

   When you save, the changes take effect immediately for the system database and all the tenant databases.

Related Information

Backup Configuration Settings [page 594]
Display the Backup Configuration Settings [page 592]
Authorizations Needed for Backup and Recovery
Log Modes
SAP HANA Backup Encryption
Schedule Backups [page 611]

11.2.1 Backup Configuration Settings

The backup configuration settings are described in the following sections.

Backint Settings

The parameter options for third-party backup tools are only available if the `Backint` agent is installed.

Backint Parameter Files

If required by the third-party backup tool, you can specify `Backint` parameter files for data backup, log backups, and the backup catalog. The content and syntax of the parameter files is tool-specific and defined by the tool vendor.

For more information, see the vendor documentation for the third-party backup tool.

> → Remember
>
> If you disable `Backint`, also check that the destination used for file-based backups is correct.

Setting	Description
Use the Same Parameter File for All	You can use the same Backint parameter file for data backups, log backups, and for backups of the backup catalog.
Data Backup, *Log Backup*, *Catalog Backup*	You can specify a different Backint parameter file for data backups, log backups, and for backups of the backup catalog.

> ℹ Note
>
> To use a parameter file, there needs to be a symbolic link pointing from `/usr/sap/<SID>/SYS/global/hdb/opt/hdbconfig/` to the actual parameter file in the directory.
>
> If a new host is added, ensure that the database services have access to the Backint agent and the parameter file.

> ℹ Note
>
> To specify parameter files for the databases in a **high isolation** system, you need to work from the system database.
>
> With high isolation, the settings are configured separately for the system database and the tenant databases. For each database, you can use a different Backint parameter file for data backups, log backups, and for backups of the backup catalog.
>
> To ensure high isolation in an SAP HANA database with many tenant databases, many Backint parameter files may be needed.

For more information, see *Isolation Level High for Backups and Third-Party Backup Tools* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

Backint Response Timeout

Setting	Description
Set Timeout (Catalog and Log Backups)	Enable a timeout for the connection to the third-party backup tool. The timeout is measured from the time of the first request to the Backint for SAP HANA agent. **i Note** This timeout is reset when data is transferred.
Timeout After (Minutes)	Specify the timeout. The default timeout is 10 minutes. If you do not specify a timeout, the default value is used. You can change the setting in increments of 5 minutes up to a maximum of 30 minutes. For more information, see *Timeout for Log Backups (Backint)* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*. 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. For more information, see *backint.log* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

Catalog Settings

Setting	Description
Destination Type	The destination type can be the file system or a third-party backup tool.

Setting	Description
Location	**For file system backups:** By default, the backup catalog is backed up to the same location as the log backups: `$(DIR_INSTANCE)/backup/log` You can specify an alternative location for backups of the backup catalog. **For Backint:** The backup data is written through the third-party backup tool. You cannot change the location. For more information, see *Destination for Backups of the Backup Catalog* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

Log Settings

Log Mode

Setting	Description
Log Mode	SAP HANA uses two log modes: `normal` and `overwrite` By default, SAP HANA runs in log mode `normal`. → Tip If you change the log mode from `overwrite` – where log backups are not written – to log mode `normal`, **you must create a full data backup** to ensure that log backups are written again, and that the database can be recovered to the most recent point in time. For more information, see *Change Log Modes* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

Log Backup

Setting	Description
Create Log Backups	Enable or disable automatic log backups. If *Create Log Backups* is disabled, the other settings in the parameter group cannot be changed. To enable log backups, the log mode must be set to `normal`. In log mode `overwrite`, you cannot change the log backup settings. For more information, see *Log Modes* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*. **⚠ Caution** During normal system operation (log mode `normal`), it is strongly recommended that you enable automatic log backups. When log segments are backed up, the space they occupied in the log area can be freed. SAP HANA can overwrite the newly freed space in the log area with new log entries. In this way, automatic log backups can prevent the log area from filling. If automatic log backups are disabled, the log area grows until the file system is full. If the file system is full, and no more log segments can be created, the database freezes.
Use Consolidated Backups	To improve the performance of log backups, SAP HANA can write multiple log segments of a service to a single consolidated log backup. If you do not use consolidated log backups, each log segment is backed up to its own backup. For more information, see *Consolidated Log Backups* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.
Maximum Size	You can configure the maximum size of consolidated log backups in increments of 8 GB. The default value is 16 GB. This means that one backup operation creates consolidated log backups with a maximum size of 16 GB. The minimum size is 8 GB. The maximum size allowed is 64 GB. For more information, see *Consolidated Log Backups* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

Setting	Description
Destination Type	The destination type can be the file system or a third-party backup tool.
Location	**For file system backups:** By default, SAP HANA log segments in the log area are backed up to: `$(DIR_INSTANCE)/backup/log` You can specify an alternative location. **For Backint:** The backup data is written through a third-party backup tool. You cannot change the location.
Back Up Logs	Specify when to back up the log segments. **At latest after specified time limit**: If log segments become full, they are backed up immediately, even if the log backup time limit has not been reached. (This is the same as log backup interval mode: immediate) **Only after specified time limit**: Log segments are backed up after the time limit you specify. This means that log segments are **not automatically backed** up if the log segments become full. (This is the same as log backup interval mode: service) For more information, see *Set the Interval Mode for Log Backups*. in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*
Time Limit	Specify the time limit for log backups. By default, the time limit is 15 minutes. The maximum permitted time limit is 1440 minutes (24 hours).

Data Backup Settings

Data Backup

Setting	Description
Destination Type	The destination type can be the file system or a third-party backup tool (Backint).
Location	**File system backups**: By default, the backups are written to the following location: `$(DIR_INSTANCE)/backup/data` You can specify an alternative location to which to write data backups. **Backint**: The backup data is written through a third-party backup tool. You cannot change the location. For more information, see *Parameters for Data Backup Settings* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.
Limit Maximum Size (File System Backups)	For file system backups, you can specify a maximum file size. The maximum file size applies to the data backups of all services. If the size of a data backup file for a service exceeds the specified limit, SAP HANA splits the file into multiple smaller files.
Maximum Size (GB)	You can specify the maximum file size of data backups in increments of 50 GB up to 2000 GB. The actual size of data backups may be smaller than the specified maximum size.

Setting	Description
Enable Multistreaming of Backint Backups	When it creates a data backup, a third-party backup tool can use multiple streams in parallel for each service. Default: 1 stream (Multistreaming is disabled) To enable multistreaming for third-party backup tools, set the number of streams to 2 or more. You can use up to 32 parallel streams. When multiple streams are used, all the parts of a backup are approximately the same size. *Threshold Triggering Multistreaming (GB)*: When multi-streaming is enabled, you can specify a threshold backup size. If a backup size is greater than the threshold, the backup uses multiple streams. If a backup is less than the specified size, only one stream is used. **i Note** To create multistreamed data backups, the third-party backup tool must also be configured to use multiple streams with good performance. For more information about the configuration of the backup tool, consult the vendor documentation. For more information, see *Configure Multistreaming with Third-Party Backup Tools* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

Data Backup Scheduler

To be able to create scheduled data backups, the data backup scheduler must be enabled. (If the scheduler is not enabled, you can still schedule backups, but the schedule will not run and they will not be created.)

Setting	Description
Enable Scheduler for Tenant Databases	Enable or disable the backup scheduler for all tenant databases. **i Note** A backup of a tenant database can be scheduled through the system database. If the assigned tenant database user is permitted, a backup of a tenant database can be also scheduled through the tenant database itself. For more information, see *Restrictions for Tenant Database Users*.
Enable Scheduler for the SYSTEMDB	Enable or disable the backup scheduler for the system database.
Assigned Database User	The user that activates the backup schedules is also used to execute the backups. If the backup scheduler is running, the associated user is displayed.

Retention Policy

You can configure the settings for the system database and each tenant database.

Setting	Description
Delete Backup Generations Automatically	Automatically schedule jobs to delete backup generations.
Retain Backup Generations Younger Than	Specify a number of days. Backup generations will be retained for at least that number of days.
Minimum No. of Retained Backup Generations	You can retain between 1 and 14 backup generations. By default, two backup generations are retained.
Options for Backup Deletion	You can delete the records of the unwanted data backup(s) from the backup catalog only, or you can delete both the records in the backup catalog and the physical backups from the file system or third-party backup tool, if you are using one.

Setting	Description
Start Daily Automatic Deletion (UTC)	You can define a point in time at which to begin automatically deleting unwanted backups. Alternatively, you can configure SAP HANA to start deleting unwanted backups at any time of the day. Alternatively, you can ask SAP HANA cockpit to select a random start time at which to delete backups every day. The period you specify is interpreted as UTC.
Assigned Database User	The user that activates the retention policy schedules is also used to perform retention actions. If the retention policy scheduler is running, the associated user is displayed.

Restrictions for Tenant Database Users

Setting	Description
Users Can Create Backups	By default, no restrictions are enforced. Tenant database users can create backups. You can prevent all users of a tenant database from creating backups. For more information about authorizations, see *Authorizations Needed for Backup and Recovery* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.
Users Have Free Choice of Backup Destination	**Note** For file-system backups only. For tenant database users with authorization to create backups. You can permit tenant database users to create backups: • In any directory • Only in the default backup destination or a subpath of the default backup destination. Users can create new subdirectories for backups below the default backup destination. **Note** Changes take effect immediately.

Related Information

Working with Third-Party Backup Tools
Isolation Level High for Backups and Third-Party Backup Tools
Timeout for Log Backups (Backint)
Delta Backups and Third-Party Backup Tools
Parameters for Backing Up the Backup Catalog
Destination for Backups of the Backup Catalog
Accumulated Backups of the Backup Catalog
Log Modes
Change Log Modes
Consolidated Log Backups
Set the Interval Mode for Log Backups
Parameters for Data Backup Settings
Schedule Backups [page 611]
Authorizations Needed for Backup and Recovery
backint.log

11.3 Create Data Backups

Using SAP HANA cockpit, you can create complete data backups and delta backups (differential backups and incremental backups).

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Procedure

1. On the *Database Overview*, with the *Administration* or *All* view and click the *Database Backups* card.

   The backup catalog is displayed.
2. (Optional) You can display more information or delete complete data backups:

To...	Steps
Display more details about a backup.	Click its row.

To...	Steps
Display an overview of backup generations.	Choose the chart icon. The chart shows the start time of the backup generation, its total size, and the sizes of the full data backup and the associated delta backups and log backups that can be used in combination for a recovery. From the overview of backup generations, you can delete backup generations. For more information, see *Delete Backup Generations*.
Change the order in which the columns are displayed.	Choose *Settings*, and use the arrow buttons. In the same way, you can also customize the backup details pages for each database.
Delete a complete data backup.	Scroll down to the row and choose *Delete*. You can remove the backup from the catalog only, or also physically delete the backup.

3. (Optional) To restrict the information displayed, choose *Filter*.

Filter By...	Description
Backup Type	Displays the backup types that were selected in the previous session.
Status	You can display backups with the status: ○ Canceled ○ Failed ○ Prepared A data snapshot has been prepared, but has not been confirmed or abandoned. ○ Running If a backup is running, click to display its progress. ○ Successful
Start Time	Display backups from a specific time range.

If you change the filter settings, your changes are retained.

4. Choose *Create Backup*.
5. Specify the backup type:

Backup Type	Description
Complete	A complete data backup includes all the data structures that are required to recover the database.

Backup Type	Description
Incremental	An incremental backup stores the data changed since the last data backup - either the last data backup or the last delta backup (incremental or differential).
Differential	Differential backups store all the data changed since the last full data backup.

The estimated size of the backup is displayed. This information is read from the system view `M_BACKUP_SIZE_ESTIMATIONS`.

For more information, see *M_BACKUP_SIZE_ESTIMATIONS System View* in the *SAP HANA SQL Reference Guide for SAP HANA Platform*.

6. Specify the backup destination type.

Option	Description
File	Writes the backup data to the file system.
Backint	Writes the backup data through a third-party backup tool. **i Note** This option is only available if a third-party backup tool is installed. The Backint parameters have no effect on the behavior of SAP HANA. For information about the Backint parameters, contact your tool vendor.

7. Specify the backup prefix.

   By default, the current date and time is proposed.

   → Tip

   It is strongly recommended to use the default prefix, as a unique timestamp makes it easier to identify archived backups.

8. Specify the backup destination.

Option	Description
For file-based backups:	Ensure that there is sufficient space at the specified backup destination. For more information, see *Estimate the Space Needed for a Data Backup* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.
For third-party backup tools:	For third-party backup tools, the destination is always: ○ System database: `/usr/sap/<SID>/SYS/global/hdb/backint/SYSTEMDB` ○ Tenant database: `/usr/sap/<SID>/SYS/global/hdb/backint/DB_<tenant_database_name>` You can only change the backup prefix.

9. To start the backup, choose *Back Up*.

   The progress of the backup is displayed.

   When all volumes have been backed up, the backup catalog overview is displayed again. Here, you can verify that the backup was completed successfully.

Related Information

Delete Backup Generations [page 623]
Estimate the Space Needed in the File System for a Data Backup
M_BACKUP_SIZE_ESTIMATIONS System View

11.4 Create a Data Snapshot (Native SQL)

You can create a data snapshot of an SAP HANA database system with one or more tenant databases. You create a data snapshot using SQL.

Prerequisites

- A data snapshot can only be created through the system database.
  It is not possible to create a data snapshot for a tenant database separately.
- The SAP HANA database (the system database and all the tenant databases) is online, and all the configured services are running.
  The system status for the system database and the tenant database is `System Running`. This is shown on the *Services* card in the *Database Overview*.

Context

A data snapshot is created in three steps that are performed in the SAP HANA database and at storage system level.

Step to Create a Data Snapshot	Description
Prepare the database for the data snapshot.	An **internal database snapshot** is created that reflects a consistent database state at the point in time it is created in the file system. **i Note** If an internal database snapshot exists, no new data backups or new data snapshots can be created. Conversely, while a data backup is running, you cannot create a data snapshot.
Create the data snapshot.	The data snapshot is created based on the previously created internal database snapshot. **i Note** To ensure its consistent state, the **data snapshot** relies on the previously created **internal database snapshot**. If the database or a database service is restarted, the **internal database snapshot** is lost. At this stage, the data snapshot is in the SAP HANA data area. To be able to create further data snapshots or data backups, you need to manually make all the content of the data area available in a separate storage location. and then confirm the data snapshot. **→ Remember** Data snapshots only offer increased data safety if they are moved or replicated to a separate storage medium. The files and directories under the mountpoint of the data area must all be stored together. The data volumes themselves must not be moved.

Step to Create a Data Snapshot	Description
Confirm or *Abandon* the data snapshot.	When the data snapshot was successfully created, you must confirm or abandon it to be able to create further data snapshots or data backups. When a data snapshot is confirmed, it is recorded in the backup catalog as successful. When a data snapshot is abandoned, it is recorded in the backup catalog as unsuccessful. The data snapshot is always **recorded in the backup catalog** - even if the internal database snapshot is lost before the data snapshot is confirmed. When the data snapshot is confirmed, you are notified whether the data snapshot can be used for a recovery. **i Note** If the confirm fails, the database snapshot is marked as unsuccessful. You should physically delete the data snapshot because it may not be possible to use it for recovery.

Procedure

To execute SQL statements, you can use the SQL console in SAP HANA cockpit. (Choose *Open SQL Console* from the *Database Overview*.)

1. Create a new internal database snapshot.

   From the system database, execute the following SQL statement:

   ```
   BACKUP DATA FOR FULL SYSTEM CREATE SNAPSHOT [COMMENT <STRING>];
   ```

   Optionally, add a comment. This comment can help to identify the data snapshot in the backup catalog.

   i Note

   `FOR FULL SYSTEM` is mandatory to create a snapshot.

   Sample Code

   ```
   BACKUP DATA FOR FULL SYSTEM CREATE SNAPSHOT COMMENT 'SNAPSHOT-2019-10-22';
   ```

   For more information, see *BACKUP DATA CREATE SNAPSHOT Statement (Backup and Recovery)* in the *SAP HANA SQL and System Views Reference*.

   An **internal database snapshot** is now created ('prepared').

> **Note**
>
> The snapshot will be created for the entire SAP HANA database, that is, the system database and all the tenant databases. It is not possible to create snapshots for an individual tenant database or only for the system database.

2. Find out the backup ID of the **internal database snapshot** in the state `PREPARED`.

> **Note**
>
> SAP HANA cannot ensure that the backup ID of the data snapshot of the system database and the backup ID of the tenant database are the same. As data snapshots are administered by the system database, you must use the backup ID of the **system database** to create the data snapshot.

Use the following SQL statement:

```
SELECT * FROM M_BACKUP_CATALOG WHERE ENTRY_TYPE_NAME = 'data snapshot';
```

> **Sample Code**
>
> ```
> SELECT BACKUP_ID, COMMENT FROM M_BACKUP_CATALOG WHERE ENTRY_TYPE_NAME =
> 'data snapshot' AND STATE_NAME = 'prepared' AND COMMENT =
> 'SNAPSHOT-2019-10-22';
> ```

Make a note of the backup ID.

> **Note**
>
> Older internal database snapshots may exist in the state `successful` or `unsuccessful`.

The database is now prepared for the data snapshot.

An **internal database snapshot** is created, reflecting a consistent database state at the point in time it is created.

> **Note**
>
> If an internal database snapshot exists, no new data backups or new data snapshots can be created.
>
> Conversely, while a data backup is running, you cannot create a data snapshot.

At this stage, all the snapshot-relevant data is only stored in the data area. To be able to use the data snapshot for a recovery later on, this data needs to be stored in a separate location.

3. In the storage system, make all the content of the data area available together in a separate storage location.

   To create the data snapshot, you can use the tool provided by your storage vendor. For more information, consult the tool documentation.

> **Note**
>
> A data snapshot contains all the persisted data in the data area. For this reason, the files and directories under the mountpoint of the data area must all be stored together.

> **→ Tip**
>
> For a recovery using a data snapshot, only the data area must be restored from the storage tool. You still can use the log area for the recovery.

> **i Note**
>
> The directory name of the data area is defined by configuration parameter `basepath_datavolumes` in the `global.ini` configuration file, in the `persistence` section.

After the data snapshot has been created in a separate storage location, it needs to be confirmed.

4. *Confirm* or *Abandon* the data snapshot.

   Use the following SQL statement:

Option	Description

Confirm	`BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID <BACKUP_ID> SUCCESSFUL <STRING>;`
	
	Confirm that the data snapshot has been successfully saved to a new storage location.
	
	You can specify an external ID to identify the data snapshot later in the storage system.
	
	Sample Code
	
	`BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID 1489592445498 SUCCESSFUL 'SNAPSHOT-2019-10-22';`

Abandon	`BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID <BACKUP_ID> UNSUCCESSFUL [<STRING>];`
	
	If the data snapshot cannot be created, or if confirmation fails, choose *Abandon*.
	
	Optionally, you can add a comment to explain why the data snapshot was not successful.
	
	Sample Code
	
	`BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID 1489592445498 UNSUCCESSFUL 'SNAPSHOT-2019-10-22 FAILED';`

For more information, see *BACKUP DATA CLOSE SNAPSHOT Statement (Backup and Recovery)* in the *SAP HANA SQL and System Views Reference*.

> **→ Tip**
>
> It is strongly recommended to confirm or abandon a data snapshot **as soon as possible after it has been created**.
>
> While the data snapshot is being prepared or created, the snapshot-relevant data is frozen. While the snapshot-relevant data remains frozen, changes can still be made in the database. Such changes will not cause the frozen snapshot-relevant data to be changed. Instead, the changes are written to

positions in the data area that are separate from the data snapshot. Changes are also written to the log.

However, the longer the snapshot-relevant data is kept frozen, the more the data volume can grow.

> **i** Note
>
> If the database or an individual database service is restarted, the **internal database snapshot** is lost. If the database snapshot is lost before the data snapshot is confirmed, the data snapshot is still written. During confirmation, the database notifies you that the data snapshot cannot be used.

After you have confirmed or abandoned a data snapshot, it is recorded in the backup catalog as either successful or unsuccessful.

> **i** Note
>
> A data snapshot now exists for both the system database and the tenant database.

The internal database snapshot that was used to create the data snapshot is discarded.

It is now possible to create further data snapshots or data backups.

Related Information

Data Snapshots
BACKUP DATA CREATE SNAPSHOT Statement (Backup and Recovery)
BACKUP DATA CLOSE SNAPSHOT Statement (Backup and Recovery)

11.5 Schedule Backups

Using SAP HANA cockpit, you can schedule data backups or delta backups to run without supervision at specific intervals. You can schedule individual backups or series of recurring backups.

Prerequisites

- The data backup scheduler must be enabled.
  The data backup scheduler is normally enabled automatically. Under certain circumstances, you may need to manually enable the data backup scheduler, which requires the XS classic engine to be running.
  For more information, see *Manually Enabling the Data Backup Scheduler*.

- You require the following authorizations:

Authorization	Purpose
`BACKUP ADMIN`	Schedule backups for the current database (system database or a tenant database). **→ Tip** To schedule backups, you can set up a user on each tenant database with the `BACKUP ADMIN` authorization or a user in the system database with the required privileges.
One of the following privileges: ○ `DATABASE BACKUP OPERATOR` ○ `DATABASE BACKUP ADMIN` ○ `DATABASE ADMIN`	Schedule backups for tenant databases through the system database. (Not backups of the system database itself) To schedule backups for a tenant database through the system database, you require a minimum SAP HANA database revision. For more information, see *SAP Note 2699762 (Backup and Recovery: Software Requirements for Scheduling Backups in SAP HANA Cockpit)*.
`SELECT`, `UPDATE`, and `DELETE` privileges for the following tables:	○ `_SYS_XS.JOB_SCHEDULES` ○ `_SYS_XS.JOBS`

For more information, see *Authorizations Needed for Backup and Recovery* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

- Backup schedules are enabled globally.
  If the data backup scheduler is not enabled, SAP HANA cockpit displays a warning, and a link to *Backup Configuration* page where you can enable the data backup scheduler.
  For more information, see *Changing the Backup Configuration Settings*.
  Under certain circumstances, it may be necessary to manually enable the data backup scheduler.
  For more information, see *Manually Enable the Data Backup Scheduler*.

  i Note
  The user that activates the backup schedules is also used to create the backups.

⚠ Caution

Backup schedules created with SAP HANA cockpit 1.0 are not compatible with SAP HANA cockpit 2.0.

Before you upgrade from SAP HANA 1.0 to SAP HANA 2.0, you must use SAP HANA cockpit 1.0 to delete all the backup schedules created with SAP HANA 1.0.

After the upgrade to SAP HANA 2.0, you need to create new backup schedules.

From SAP HANA cockpit 2.0, to schedule backups for an SAP HANA 1.0 database, you must be logged onto that database. You cannot schedule backups for SAP HANA 1.0 databases through the system database.

Procedure

1. From the *Database Directory*, select a database:

To Schedule Backups For...	Perform the Following Steps...
A tenant database **(through the system database)**	1. Log onto the **system database**. The *Database Overview* page is displayed. 2. From the *Database Overview*, choose *Database Management*. 3. Choose *Backup Schedules*. The scheduling calendar is displayed.
The current database	1. Log onto the database for which you want to schedule backups (the system database or a tenant database). 2. From the *Database Backups* card, choose *Backup Schedules*. The scheduling calendar is displayed for the current database only.

2. To create a new backup schedule, choose *+ (Create Schedule)*.

   Alternatively, click the calendar on the day for which you want to create the schedule.
3. Specify whether to schedule a series of recurring backups or a one-off backup.
4. If you are working through the system database, select a database for which to create the schedule.
5. If you are scheduling a series of backups, specify a name for the schedule.

   It is recommended to choose a schedule name that enables you to easily identify the schedule.
6. Specify the backup type.

Option	Description
Complete	A complete data backup includes all the data that is required to recover the database to a consistent state.
Differential	Backs up all the data changed since the last full data backup (complete data backup or data snapshot).
Incremental	Backs up the data changed since the last full data backup (complete data backup or data snapshot) or the last incremental or differential backup.

> **i Note**
>
> Currently, scheduling data snapshots is not supported.

7. Specify the *Destination Type*.

Option	Description
File	To write the backups to the file system, select *File*.

Option	Description
	If necessary, you can specify a new destination or change the default destination. For more information, see *Parameters for Data Backup Settings*.
Backint	This option is only displayed if you are working with a third-party backup tool. To create backups using a third-party backup tool, select the destination type *Backint*. If needed, specify the *Backint Parameters*. For more information, see *Working with Third-Party Backup Tools*.

8. Specify a *Backup Prefix*.

> → Tip
>
> By default, the name of each scheduled backup is prefixed with the timestamp of the start of the backup. The placeholder `<[date]_[time]>` is automatically converted to the current timestamp.
>
> To be able to more easily identify archived backups, it is strongly recommended to use the default, as it provides a unique prefix for each backup.

9. Specify a *Backup Destination*.

   By default, file-based data backups are written to the following subdirectories:

   - System database: `$DIR_INSTANCE/backup/data/SYSTEMDB`
   - Tenant database: `$DIR_INSTANCE/backup/data/DB_<tenant_database>`

   For file-based data backups, you have the option to change the default backup destination.

   The default backup destination for third-party backup tools cannot be changed.

10. (Optional) Add a comment.

    A comment helps you to later identify the backups in the backup catalog.

11. If you are scheduling a series of backups, specify the recurrence pattern.

    You can schedule the backup to run each week, once each month, every two months, quarterly, twice a year, or once a year.

12. Specify the recurrence options for the backup schedule.

    The recurrence options depend on whether you specified a weekly or month-based recurrence pattern.

Recurrence Setting	Description
Time Zone	Select the time zone in which you want to specify the backup time. You can select any time zone that is convenient for you.
Create Backups At	Specify a time to run the scheduled backups in the selected time zone. You can specify a time manually or choose the clock icon to select a time from the list.

Recurrence Setting	Description
(Weekly) *Create Backups On*	Specify on which days of the week you want the backups to be created. You can select one or more days.
(Month-based) *Create Backups On (UTC)*	Specify a day or a week of the month. *Day of Month:* You can select the first day, the 15th day, or the last day in the month. *Week of Month:* Specify which day of the month to create the backup. For example, the last Sunday in the month or the first Monday in the month.
Activate Schedule On	Specify a day in the selected time zone. The first backup scheduled will be created on this day or as soon as possible afterwards when the schedule criteria are fulfilled.

When you have specified the recurrence options, SAP HANA cockpit displays when the first backup will be created. The time is shown in both the time in the specified time zone and UTC.

If you specified a month-based recurrence pattern, the months in which the backups will be created are also shown. For example, October or Monthly.

> **i Note**
>
> All times specified are interpreted as UTC. If the local time changes due to daylight saving time, the UTC time for the scheduled backup does not change.

13. Choose *Review*.

    A summary of the schedule options is displayed.

    To make changes, choose *Edit* for an option group.
14. When the schedule settings are complete, choose *Save Schedule*.

    The new schedule is saved and is active immediately.

    The next backup scheduled will be created at the time displayed.

> **⚠ Caution**
>
> If SAP HANA is **offline** at a time for which backups are scheduled, scheduled backups will not run.
>
> When SAP HANA is running again, skipped backups are **not automatically rescheduled**.

Related Information

Authorizations Needed for Backup and Recovery

Manage Backup Schedules [page 619]
Manually Enable the Data Backup Scheduler [page 616]
Enable the XS Job Scheduler (XS Classic) [page 617]
Enable the XS Engine (XS Classic) [page 618]
SAP HANA Backup Types
Parameters for Data Backup Settings
Working with Third-Party Backup Tools
SAP Note 2699762
Change the Backup Configuration Settings [page 592]

11.5.1 Manually Enable the Data Backup Scheduler

When backup schedules are created, they are activated automatically. Under certain circumstances, you may need to manually enable the data backup scheduler, which requires the XS classic engine to be running.

Prerequisites

You've navigated to the *Database Overview* page of the database you want to manage. See *Getting to the Database Overview Page* in the *SAP HANA Administration with SAP HANA Cockpit* guide.

Context

> **i Note**
>
> The user that activates the backup schedules is also used to execute the backups.

Procedure

1. On the *Database Overview*, with the *Administration* or *All* view selected, click the *Backup Configuration* link on the *Database Backups* card.
2. Go to *Data Backup Settings* > *Data Backup Scheduler*.
3. Choose *Edit*.
4. Set Enable Data Backup Scheduler to YES.
5. *Save*.

   The data backup scheduler and the XS Job Scheduler are enabled.

> **i Note**
>
> If it is not possible to enable the data backup scheduler, you may need to manually enable the XS Job Scheduler, and possibly also the XS Engine.
>
> For more information, see *Enable the XS Job Scheduler (XS Classic)* and *Enable the XS Engine (XS Classic)*.

Related Information

Enable the XS Job Scheduler (XS Classic) [page 617]
Enable the XS Engine (XS Classic) [page 618]
Schedule Backups [page 611]

11.5.1.1 Enable the XS Job Scheduler (XS Classic)

To schedule data backups or delta backups using SAP HANA cockpit, the XS Job Scheduler (XS Classic) must be enabled.

Prerequisites

The XS Job Scheduler requires the XS Engine to be enabled.

Context

Normally, the XS Engine is enabled by default. In some cases, you may need to activate the XS Engine manually.

The XS Job Scheduler is enabled separately for the system database and for each tenant database.

Procedure

1. For the **system database**, the XS Job Scheduler must be enabled in the `nameserver.ini` file.

   To enable the XS Job Scheduler, use the following SQL statement:

   ```
   ALTER SYSTEM ALTER configuration ('nameserver.ini','SYSTEM') SET
   ('scheduler','enabled')= 'true' WITH reconfigure;
   ```

2. For each **tenant database**, the XS Job Scheduler must be enabled in the `xsengine.ini` file.

   To enable the XS Job Scheduler, use the following SQL statement:

   ```
   ALTER SYSTEM ALTER configuration ('xsengine.ini','SYSTEM') SET
   ('scheduler','enabled')= 'true' WITH reconfigure;
   ```

Related Information

Schedule Backups [page 611]
Enable the XS Engine (XS Classic) [page 618]

11.5.1.2 Enable the XS Engine (XS Classic)

To schedule backups using SAP HANA cockpit, the XS scheduler must be enabled. The XS scheduler requires the XS engine to be running.

Context

Normally, the XS engine is enabled by default in an SAP HANA system. In some cases, you may need to activate the XS engine manually.

The XS engine is enabled separately for the system database and for each tenant database.

Procedure

1. For the **system database**, the XS engine must be enabled in the `nameserver.ini` file.

   To enable the XS engine, you can use the following SQL statements:

   ```
   ALTER SYSTEM ALTER configuration ('nameserver.ini','SYSTEM') SET
   ('httpserver','embedded')= 'true' WITH reconfigure;
   ```

   ```
   ALTER SYSTEM ALTER configuration ('nameserver.ini','SYSTEM') SET
   ('httpserver','workerpoolsize')= '5' WITH reconfigure;
   ```

   > → Tip
   >
   > It is recommended that you set the workerpool size to 5 for the system database. However, if you need to schedule many backup jobs, consider increasing the value in accordance with your system requirements.

2. For each **tenant database**, the XS engine must be enabled in the `xsengine.ini` file.

To enable the XS engine, you can use the following SQL statement:

```
ALTER SYSTEM ALTER configuration ('xsengine.ini','SYSTEM') SET
('httpserver','embedded')= 'true' WITH reconfigure;

ALTER SYSTEM ALTER configuration ('xsengine.ini','SYSTEM') SET
('httpserver','workerpoolsize')= '5' WITH reconfigure;
```

> → Tip
>
> It is recommended that you set the workerpool size to 5 for a tenant database. However, if you need to schedule many backup jobs, consider increasing the value in accordance with your system requirements.

Related Information

Schedule Backups [page 611]
Enable the XS Job Scheduler (XS Classic) [page 617]

11.5.2 Manage Backup Schedules

Using SAP HANA cockpit, you can display an overview of backup schedules, pause or reactivate schedules, and delete schedules.

Procedure

1. From the system database, on the *Database Overview*, choose *Database Management*.

   A status overview of the system database and the tenant databases is displayed.
2. Choose *Backup Schedules*.

   An overview of the backups scheduled for the system database and the tenant databases is displayed.
   a. You can display the schedules for a week or for a whole month. Choose *1 Month* or *1 Week* from the pull-down menu.

      If a schedule is set to be executed at a future date, scroll forward in the schedule calendar to see that schedule.
   b. Choose a task.

To...	Steps
Display the details of a backup schedule.	Click the backup in the calendar.
Change a backup schedule.	1. Select a backup schedule and choose *More Details*.

To...	Steps
	2. Choose *Edit*. 3. Make the desired changes. > **i Note** > It is not possible to change the recurrence of a schedule between weekly and monthly. For a schedule with weekly recurrence, you can change the days of the week on which the schedule is executed. 4. Choose *Save*.
Create a new backup schedule.	Choose *+ (Create Schedule)* or click the date in the calendar. For more information, see *Schedule Backups*.
Pause or reactivate a backup schedule.	1. Select a backup schedule and choose *More Details*. 2. Choose *Pause*. To reactivate a paused schedule, choose *Activate*.
Delete a backup schedule.	1. Select a backup schedule and choose *More Details*. 2. Choose *Delete* The backup schedule is deleted permanently.

Related Information

Schedule Backups [page 611]

11.6 Cancel a Backup

You can cancel a running data backup or a delta backup (differential or incremental).

Context

The option to cancel a backup is only available while the backup is running.

> **i Note**
>
> In some situations, it may **not** be possible to cancel a running backup. For example, if it is not possible to access internal locks, or if a file cannot be written to an NFS mount.

Procedure

1. From the backup dialog, choose *Cancel Backup*.
2. Confirm your decision.

   The backup is canceled and you are notified.

Results

After you have canceled a backup, you can start a new data backup.

> → Tip
>
> If you canceled a running backup performed by a **third-party backup tool**, it is recommended to ensure that any incomplete backups are physically deleted.

11.7 Housekeeping: Deleting and Archiving Backups

It is recommended to regularly check whether old full backups or backup generations can be deleted.

You can delete full backups and backup generations if they are no longer needed for a recovery, or to keep your backup storage space at an optimum level.

For more information, see *Deleting Backup Generations*.

Archiving Backups

Full backups that need to be retained for an extended period can be archived in a secure location and then removed from the backup catalog. Ensure that archived backups cannot be accessed directly by SAP HANA and cannot be deleted.

> → Tip
>
> It is important to regularly truncate the backup catalog because, as it increases in size, it can consume a lot of storage space and also take longer to write each new backup.

An archived **full data backup** can still be used to recover SAP HANA, even if it is not recorded in the backup catalog.

If you need to ensure that you can recover SAP HANA from older log backups and delta backups, you need to retain backups of the backup catalog.

Related Information

Prerequisites for Deleting Old Backups [page 622]
Delete Backup Generations [page 623]
Delete Backups [page 625]
Backup Catalog

11.7.1 Prerequisites for Deleting Old Backups

Before you delete old backup generations or individual full backups, some prerequisites must be met.

- You have the appropriate system privilege:

BACKUP ADMIN	Connect directly to a system database or a tenant database and delete backup and backup generations of that database. For more information, see *BACKUP CATALOG DELETE Statement (Backup and Recovery)*.
DATABASE ADMIN	Connect to a system database and delete backup generations or individual full backups of a tenant database.

- To delete backups and backup generations of a tenant database, the tenant database must be online.
- To delete backup generations, you need to decide from which time onwards you want to retain data backups.
- Before you physically delete backup generations or individual full backups, ensure that the backups retained are accessible and consistent, so that, if needed, they can be used to recover the database.
- At least one full data backup must remain in the backup catalog.

Deleting Data Snapshots

When you delete a data snapshot from SAP HANA cockpit, **only its record is removed from the backup catalog**.

SAP HANA does not physically delete a data snapshot. To physically delete a data snapshot, you must delete it manually in the repository where it is stored.

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.

Related Information

Delete Backups [page 625]
Rebuilding the Backup Catalog
Data Snapshots

11.7.2 Delete Backup Generations

Using SAP HANA cockpit, you can delete backup generations that are no longer needed for a recovery.

Context

You can delete the backup catalog records of backup generations, but retain the associated physical backups. Optionally, you can delete the records from the backup catalog and also the associated physical backups.

> **i** Note
>
> You can automate the deletion of backup generations by configuring the retention policy.
>
> For more information, see *Change the Backup Configuration Settings*.

A backup generation comprises the following backups:

A Backup Generation Consists Of...	...Backups
Successfully created full backups:	• Complete data backup **OR** • Data snapshot
AND Backups that were created after the full backup and up to the start time of the next successful full backup:	• Delta backups (differential or incremental backups) • Log backups • Backups of the backup catalog

Procedure

1. From the *Database Directory*, select a database.

Option	Description
To access a system database	From the *Database Overview* of the system database, locate the *Database Backups* card and click it. An overview of information from the backup catalog is displayed for the system database.
To access a tenant database from its system database	From the *Database Overview* of the system database: 1. Choose *Database Management*. 2. Click the backup entry. The backup catalog overview for the tenant database is displayed.

Option	Description
To access a tenant database directly	From the *Database Overview* of the tenant database, locate the *Database Backups* card and click it. An overview of information from the backup catalog is displayed for the tenant database. **i Note** To delete backups directly from a tenant database, you need the BACKUP ADMIN privilege for the tenant database, and the *Restrictions for Tenant Database Users* permit the tenant database user to make changes. For more information, see *Authorizations Needed for Backup and Recovery* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)* and *Backup Configuration Settings*.

2. To display an overview of backup generations, choose the chart icon.

   The chart shows the start time and end time of the backup generation, its total size, and the sizes of the full data backup and the associated delta backups, log backups, and backups of the backup catalog that can be used for a recovery.

   i Note

   The end time of a backup generation is the start time of the next backup generation.

   To customize the information displayed, choose *Filter*.

   In the dialog box, you can specify a time range from which to display the backup generations.
3. Click the row of the most recent backup generation that you want to retain.

   A summary of information about the backup generation is displayed.
4. To delete one or more backup generations, choose *Delete Backup Generations*.

   ⚠ Caution

   When you select one backup generation, **all the backup generations** that are **older than the selected backup generation** will be deleted.

   This includes the backup generations that are displayed, and also the backup generations that are outside of the time filter, and therefore not currently displayed.
5. In the dialog box, select an option:

Option	Description
Remove from backup catalog only	Delete only the backup catalog records of backup generations, but retain the associated physical backups.
Also delete physically	Delete the records from the backup catalog and also the associated physical backups. You can delete the physical backups in either the file system or a third-party backup tool, or both. **⚠ Caution** When you confirm, the records of the backups are deleted from the backup catalog immediately, even though it may take some more time for the physical backups to be deleted. **i Note** SAP HANA can only physically delete backups that are in the location recorded in the backup catalog. SAP HANA cannot physically delete backups that have been moved to a different location.

6. To confirm, choose *Delete*.

   The backup generations are deleted in accordance with the options you specified.

Related Information

Change the Backup Configuration Settings [page 592]

11.7.3 Delete Backups

Using SAP HANA cockpit, you can display an overview of complete data backups (and delta backups), and also delete individual full data backups.

Context

i Note

To ensure that no backups are deleted that would prevent SAP HANA from being recovered, it is not possible to delete delta backups and log backups individually.

You can delete the backup catalog records of individual full backups, but retain the associated physical backups. Optionally, you can delete the records in the backup catalog and also the associated physical backups.

For example, to comply with legal requirements for data retention, you may wish to retain specific historical data backups, but without the intention of using them for a production database recovery.

> **i Note**
>
> If a full backup is physically available, but not recorded in the backup catalog, that full backup can still be used to recover the database.
>
> To be able to recover SAP HANA using **third-party data backups** that are not recorded in the backup catalog, the data backups must have a unique prefix.

Procedure

1. From the *Database Directory*, select a database.

Option	Description
To access a system database	From the *Database Overview* of the system database, locate the *Database Backups* card and click its title. An overview of information from the backup catalog is displayed for the system database.
To access a tenant database from its system database	1. From the *Database Overview* of the system database, choose *Database Management*. 2. Select the tenant database. The backup catalog overview for the tenant database is displayed.
To access a tenant database directly	From the *Database Overview* of the tenant database: 1. Choose *Database Management*. 2. From the overview of the databases in the SAP HANA system, select the tenant database. The backup catalog overview for the tenant database is displayed. **i Note** To delete backups directly from a tenant database, you need the BACKUP ADMIN privilege for the tenant database, and the *Restrictions for Tenant Database Users* permit the tenant database user to make changes. For more information, see *Authorizations Needed for Backup and Recovery* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)* and *Backup Configuration Settings*.

   a. To customize the information displayed, choose *Filter*.

From the dialog box, you can filter the following information:

Backup Type	By default, *Complete Data Backup* and *Data Snapshot* are selected. To display information about delta backups, select *Differential Backup* and *Incremental Backup* and confirm.
Status	You can display backups with the status: ○ Canceled ○ Failed ○ Prepared A data snapshot has been prepared, but has not been confirmed or abandoned. **i Note** To create a data snapshot, you need to use native SQL. ○ Running ○ Successful
Start Time	You can display backups from a specific time range.

b. To change the order in which the columns are displayed, choose *Settings*, and use the arrow buttons.

2. To delete a full backup (complete data backup or data snapshot), choose the *Delete* icon for that backup.

In the dialog box, specify whether to:

Remove from backup catalog only	Delete only the backup catalog record of a full backup, but retain the associated physical backup.
Also delete physically	Remove the record of the backup from the backup catalog and delete the associated physical backup. You can delete the physical full backup in either the file system or a third-party backup tool. **⚠ Caution** When you confirm, the record of the backup is deleted from the backup catalog immediately, even though it may take some more time for the physical backup to be deleted. **i Note** SAP HANA can only physically delete backups that are in the location recorded in the backup catalog.

SAP HANA cannot physically delete backups that have been moved to a different location.

3. To confirm, choose *Delete*.

   Before a backup is physically deleted, the following plausibility checks are performed:

   - **For a file-based backup:** The system checks the backup ID.
   - **For third-party backup tools:** The system checks the external backup ID (EBID) and whether the path to the backup is identical to the backup location of the current database.

   If the plausibility check is successful, the system starts deleting the physical backup in the background.

   > **i Note**
   >
   > The delete operation continues until all the parts of the selected backup have been deleted.
   >
   > If the system or a service is stopped and restarted, the delete operation is automatically resumed.

   You can monitor the progress of the deletion operation in the `backup.log` file.

Related Information

Delete Backup Generations [page 623]
Create Data Backups [page 603]
Backup Configuration Settings [page 594]

11.8 Recover a Database

Using SAP HANA cockpit, you can recover an SAP HANA database.

Context

A tenant database is recovered through its system database.

A tenant database can be recovered to its most recent state or to a specific point in time.

Using SAP HANA cockpit, a system database can only be recovered to its most recent state.

> **i Note**
>
> A recovery to the most recent state or to a point in time is equivalent to the SQL statement `RECOVER DATABASE` (not `RECOVER DATA`).

For more information, see *RECOVER DATABASE Statement (Backup and Recovery)* and *RECOVER DATA Statement (Backup and Recovery)* in the *SAP HANA SQL Reference Guide*.

Procedure

1. From the *Database Directory*, select the system database.

   The *Database Overview* is displayed.

Option	Description
To recover the system database	From the *Database Backups* card, choose *Recover database*. If the system database is still running, you are prompted to shut it down. If the system database is shut down for recovery, all its tenant databases are automatically shut down as well. The whole SAP HANA system is not available until the recovery of the system database has been completed.
To recover a tenant database	Choose *Database Management*. Select the row with the tenant database and choose *Recover*. If the database is not already shut down, you are prompted to shut it down.

   Follow the steps on the screen.

2. When the database is shut down, specify the state to which you want to recover the database.

Option	Description
Recover to the most recent state	Recover the database to a state as close as possible to the current time. → Tip Using the most recent complete data backup makes for a faster recovery.
Recover to a specific point in time	Specify a time zone and a point in time to which to recover the tenant database. i Note Any changes that were made **after** the specified point in time will not be in the recovered tenant database. i Note If you specify a point in time in the future, the effect will be the same as recovering the database to the most recent state.

3. Specify the location of the most recent backup catalog.

Option	Description
Backint location only	If a third-party backup tool is selected, Backint is searched.
Default location	For file system backups, the location for the backups of the backup catalog is defined using the parameter `basepath_catalogbackup`. The default setting for `basepath_catalogbackup` is: `$DIR_INSTANCE/backup/log` By default, log backups for tenant databases are written to a tenant-specific subdirectory. By default, backups of the backup catalog are written to the same tenant-specific subdirectory as the log backups.
Alternative location	If the backup catalog is not in the default location, specify its location. For more information, see *Destination for Backups of the Backup Catalog* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

An overview of available backups is displayed.

4. Select the data backup to use for the recovery.
5. Specify whether to use delta backups (differential and incremental backups).

> ⚠ Caution
>
> The complete data backups and the delta backups must be in the same location for the recovery to work correctly.

For more information about delta backups, see *Delta Backups*.

6. If necessary, you can change the location to be searched for data backups and delta backups.

   To specify additional locations for log backups, choose *Add more*.

   If you leave the locations empty, SAP HANA uses the backup locations specified in the backup catalog.
7. Check whether the backups are available.

   Here, you can decide whether to check if all the backups needed are available and can be accessed **before the recovery starts**. The availability check is performed at the beginning of the recovery.

> ℹ Note
>
> SAP HANA does not check the integrity of the backups content on block level.
>
> For more information, see *Manually Checking Whether a Recovery is Possible* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

8. Specify whether to initialize the log area.

> ⚠ Caution
>
> If you initialize the log area, the content of the log area is lost.

> No log records from the log area can then be replayed during the recovery. Only the log backups can be used.

In the following situations, you **must** initialize the log area:

- The log area is unusable.
- You are recovering the database to a different system.

9. Choose *Review*.

   An overview of the settings for the recovery is displayed.

   To change any settings, choose *Edit*.

   All the settings that you specified are retained until you change them.

10. To display the SQL statement to be used for the recovery, choose *Display SQL Statement*.

    For more information, see *Recovering a Database Using Native SQL* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

11. To perform the recovery, choose *Start Recovery*.

    The progress of the recovery for each SAP HANA service is displayed.

Results

When the recovery is completed, a message confirms this, and shows the point in time to which the database was recovered.

> **i Note**
>
> The SQL statement used for a recovery is recorded in `backup.log`. For a point-in-time recovery, the point in time is specified in the SQL statement as **UTC**.
>
> The time at which the recovery was started and completed is recorded in `backup.log` as **local server time**, not UTC.
>
> For more information, see *Diagnosis Files for Backup and Recovery* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

> **i Note**
>
> The point in time that SAP HANA returns after a recovery may be **before** the point in time that you specified for the recovery. This is because the point in time that was actually reached in the recovery is that of the most recent global COMMIT to the database that was recovered.

> **i Note**
>
> When you recover and restart a system database, its tenant databases are not automatically restarted. You should first check that the system database was recovered successfully, then restart the tenant databases manually.

The SAP HANA database is now online and can be used by applications.

Related Information

Prerequisites for Database Recovery
Destination for Backups of the Backup Catalog
Delta Backups
Manually Checking Whether a Recovery is Possible
Cancel a Recovery [page 635]
Resume a Canceled Recovery [page 637]
Diagnosis Files for Backup and Recovery
Create a Data Snapshot (Native SQL) [page 606]
Recover SAP HANA From a Data Snapshot [page 632]
RECOVER DATABASE Statement (Backup and Recovery)
RECOVER DATA Statement (Backup and Recovery)

11.8.1 Recover SAP HANA From a Data Snapshot

Using SAP HANA cockpit and storage system tools, you can use a data snapshot to recover a complete SAP HANA database.

Prerequisites

- Using a data snapshot, you can recover a complete SAP HANA database, that is, the system database and all the tenant databases.
- To recover SAP HANA from a data snapshot, the data snapshot must be available in the data area of the storage system.

  > **i** Note
  >
  > Before you make the data snapshot available in the data area of the storage system, you must first shut down the database.

- When you recover SAP HANA from a data snapshot, you need to recover **both** the system database and each tenant database separately.
  You first need to recover the system database. After the system database has been recovered successfully, each tenant database is recovered separately. The tenant databases cannot be recovered together in one single operation.
- To recover a database from a data snapshot, you can optionally also use delta backups and log backups.

Procedure

First, recover the **system database** from the data snapshot.

1. From the *Database Directory*, open the system database and shut down the database.
2. From the *Database Overview*, then choose *Stop System*.
3. Follow the on-screen instructions to shut down the database.

   The database is shut down.

   In *Manage Services*, the system status is shown as *Stopped*.
4. Go to the *Database Overview*.
5. Outside of SAP HANA cockpit, make the data snapshot (the content of the data area) available in the data area of the storage system.

   The data snapshot must be accessible by the `<SID>adm` user.
6. From SAP HANA cockpit, go to the *Database Overview*, and choose *Recover database* from the *Database Backups* card.
7. Specify the recovery target.

   You can recover the database to its most recent state or specify a point in time to which to recover the database.
8. Specify the location of the most recent backup catalog.

Option	Description
Backint location only	If a third-party backup tool is selected, Backint is searched.
Default location	For file system backups, the location for the backups of the backup catalog for the system database is defined using the parameter `basepath_catalogbackup`. The default setting for `basepath_catalogbackup` is: `$DIR_INSTANCE/backup/log` By default, log backups for tenant databases are written to a tenant-specific subdirectory. By default, backups of the backup catalog are written to the same tenant-specific subdirectory as the log backups.
Alternative file system location	If the backup catalog is not in the default location, specify its location. For more information, see *Destination for Backups of the Backup Catalog* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

   An overview of available backups is displayed.
9. Select the data snapshot.
10. Specify whether to use delta backups.

> **i Note**
>
> If you wish to recover SAP HANA using delta backups (differential or incremental backups), you must also use log backups. If log backups are not available, you can only recover using a full data backup.

11. Specify the location of the log backups, if available.
12. To ensure that a backup exists at the specified location, check the availability of the backups.

13. If necessary, initialize the log area.

    If you initialize the log area, the content of the log area is lost. No records from the log area can then be replayed. The records from the log backups are replayed if they are needed.

    > ⚠ Caution
    >
    > Disabling log backups may cause significant loss of data.

    In the following situations, you **must** initialize the log area:

    - The log area is unusable
    - You are recovering the database to a different system

14. To recover the system database, choose *Start Recovery*.

    The progress of the recovery is displayed.

    When the recovery is completed, a message confirms this, and shows the point in time to which the database was recovered.

    The system database is now online, but cannot yet be used by applications.

    You should now first check that the system database was recovered successfully, then restart the tenant databases manually.

    For more information, see *Recover all Tenant Databases From a Data Snapshot*.

Related Information

Recover all Tenant Databases From a Data Snapshot [page 634]
Recovering an SAP HANA Database
Data Snapshots
Authorizations Needed for Backup and Recovery
Destination for Backups of the Backup Catalog

11.8.1.1 Recover all Tenant Databases From a Data Snapshot

When the system database has been recovered successfully from a data snapshot, you need to recover all the tenant databases from the same data snapshot.

Prerequisites

- The system database has been successfully recovered.
- The data snapshot that was used to recover the system database remains available in the data area of the storage system.

Context

When the system database has been recovered successfully from a data snapshot, only the system database is started automatically. When you have checked that the system database was recovered correctly, you need to then recover each tenant database separately.

Procedure

1. From the *Database Directory*, select the recovered system database and go to the *Database Overview*.
2. Choose *Database Management*.

   An overview of the databases in the SAP HANA system is displayed.
3. Select a tenant database.
4. Choose *Recover Tenant*.

   Follow the on-screen instructions.
5. To recover the tenant database, choose *Start Recovery*.

   The progress of the recovery for each SAP HANA service is displayed.

   When the recovery is completed, a message confirms this, and shows the point in time to which the database was recovered.

   The SAP HANA tenant database is now online and can be used by applications.

   To recover the remaining tenant databases, repeat the steps for each tenant database.

Related Information

Recover SAP HANA From a Data Snapshot [page 632]

11.9 Cancel a Recovery

You can cancel a recovery while it is in progress. After a recovery is canceled, it needs to be repeated or resumed before work can continue in the database.

Context

While a recovery is in progress, the option to cancel it is displayed.

Procedure

1. From the recovery progress view, choose *Cancel Recovery*.
2. Confirm your decision.

Results

> ⚠ Caution
>
> After a recovery has been canceled, the database has an inconsistent state.
>
> SAP HANA automatically prevents a database with an inconsistent state from being started.
>
> For this reason, the only way to make the database available is to repeat or resume the recovery.
>
> For more information about resuming a recovery, see *Resume a Canceled Recovery* in the *SAP HANA Administration Guide (SAP HANA Database Backup and Recovery)*.

> i Note
>
> If you attempt to restart the database after a recovery is canceled, the following message is written to the nameserver trace file:
>
> `Cannot start the service 'nameserver' at '<host:SQL Port>' responsible for the volume '<volume number>' because of an error during recovery.`

Related Information

Resume a Canceled Recovery [page 637]

11.10 Resume a Canceled Recovery

Instead of repeating an entire recovery, it is possible to resume a recovery that was canceled or aborted. It is normally only necessary to resume a recovery in exceptional circumstances.

Prerequisites

> **i** Note
>
> After a recovery is canceled or aborted, the SAP HANA database cannot restart. Before work can continue in the database, the recovery must be completed.

- It is possible to resume a recovery that uses a full backup (a data backup or a data snapshot) with delta backups and log backups.
  A recovery can only be resumed after the recovery from the full backup is completed.
  If the recovery is canceled or fails during recovery from a data backup, the recovery cannot be resumed.

  > **!** Restriction
  >
  > A recovery from **only a full data backup** cannot be resumed.
  >
  > If a recovery from only a full data backup is canceled or fails, the recovery needs to be repeated from the beginning.
- To resume a recovery from a data snapshot, you can use the same data snapshot that you used when you started the recovery.
  The data snapshot does not need to be replicated to the data area again. You only need to ensure that the required delta backups and log backups are available.
- It is possible to resume a recovery both with file-based backups and with third-party backup tools.
- A resumed recovery must use the same backup catalog as the recovery that was canceled or aborted.

Disk Sizing

When you resume a recovery, disk space requirements can temporarily increase. If you resume a recovery, as a general rule, ensure that you are working with enough disk space to contain at least the data backup and the delta backups.

Context

In many situations, a recovery can be repeated from the beginning in only a short time, and the database can be running again with only minimal loss of uptime. However, in some situations, having to repeat a recovery from the beginning may cause a significant delay. In such situations, the option to resume a recovery can save a considerable amount of time, both with a very large database or a relatively small database.

Fallback Points

During a recovery, SAP HANA automatically defines fallback points, which mark the points after which it is possible to resume a recovery.

The first fallback point occurs after the recovery from the complete data backup. After recovery from a delta backup is completed, another fallback point is set. Fallback points are also written during log recovery.

A recovery can be resumed from the latest fallback point. The part of the recovery from before the fallback point does not need to be repeated.

> **i** Note
>
> If you resume a recovery to an earlier point in time than for the canceled recovery, the log fallback points are not used. In this situation, only the fallback point for the recovery from the data backup (and, if used, also the delta backups) can be used; the log recovery needs to be repeated from the beginning.

After a recovery has been successfully completed, the fallback points are invalidated. It is then no longer possible to perform a new recovery based on those fallback points.

Configuring Fallback Points

You can define how often fallback points are written during log recovery.

Configure the parameter `log_recovery_resume_point_interval` in the `global.ini` configuration file.

The time interval that you specify translates to the maximum acceptable database uptime lost while the log recovery is resumed after a recovery from the data backup (and possibly also the delta backups) has been completed.

Viewing Fallback Points

Each SAP HANA service defines its own service-specific fallback point after recovery from a data backup or delta backups. Together, these service-specific fallback points make up the fallback point for the whole database. When a fallback point has been written for **all the services** in the database, the fallback point for the whole database is recorded in `backup.log`.

> **i** Note
>
> No fallback points are written during log recovery.

> **Example**
>
> The following example shows one fallback point in the `backup.log` file.
>
> In this example, the fallback point is written to `fallback_databackup_0_1`:
>
> ```
> 2018-10-22T17:41:46+02:00 P012898 15c215ef3c7 INFO RECOVERY fallback point
> written into /usr/sap/DB1/SYS/global/hdb/data/mnt00001/hdb00002.00004/
> fallback_databackup_0_1
> ```

For more information, see *backup.log*.

Procedure

Resume a Canceled Recovery

1. Re-start the recovery.

   A partially completed recovery that can be resumed is displayed in the backup overview.

   For more information, see *Recover a Database*.
2. In the recovery dialog, select the backup with which the recovery can be resumed.
3. Follow the steps described on-screen to complete the recovery.

Results

When the recovery is complete, a message confirms this, and shows the time to which the database was recovered.

The SAP HANA database is now online and can be used by applications.

Related Information

backup.log
Recover a Database [page 628]
RECOVER DATABASE Statement (Backup and Recovery)

11.10.1 Resuming a Canceled Recovery Using Native SQL

You can resume a recovery by using the `recoverSys.py` tool and the SQL recovery command with `USING RESUME`.

For more information about `recoverSys.py`, see *Recover a System Database to a Point in Time*.

> **i** Note
>
> `USING RESUME` can only be used if a fallback point already exists.
>
> If no fallback point exists, an error is returned.

For simple SQL statements, `USING RESUME` is appended.

For SQL statements with options such as `IGNORE DELTA DATA BACKUPS` or `CHECK ACCESS`, `USING RESUME` should be placed **after the path** and **before the option**.

> Sample Code

```
RECOVER DATABASE UNTIL TIMESTAMP '2018-10-22 17:41:46' USING RESUME
```

```
RECOVER DATABASE FOR Tenant_1 UNTIL TIMESTAMP '2018-10-22 17:41:46' USING
CATALOG PATH ('/hana/DB1/backup/catalog') USING LOG PATH ('/hana/DB1/backup/
log') USING RESUME CHECK ACCESS ALL;
```

Related Information

RECOVER DATABASE Statement (Backup and Recovery)
RECOVER DATA Statement (Backup and Recovery)
Recover a System Database to a Point in Time

11.11 Copy a Database

Using SAP HANA cockpit, you can create a copy of an SAP HANA database by recovering it to the same system or to a different system.

Procedure

1. Locate the database to copy to.

   > **i** Note
   >
   > This database will be overwritten by the recovery.

Database to Copy To	Perform the Following Steps...
System database	1. From the *Database Overview* for the system database, go to the *Database Backups* card, and choose *Copy Database*. 2. Follow the instructions on the screen. The options are described in the following sections. The actual sequence of steps that you perform depends on the specific options that you choose.
Tenant database	1. From the *Database Overview*, choose *Database Management*. 2. Select the tenant database that you want to copy, and choose *Copy*. 3. Follow the instructions on the screen. The actual sequence of steps that you perform depends on the specific options that you choose.

2. Specify the database copy type.

Database Copy Type	Description
Full data backup only	Create a copy of the database from the start time of the full data backup.
Data and log backups	Create a copy of the database to a specific point in time. In the next step, you are prompted to specify the time zone, the date and the time.

3. Specify whether to use a backup catalog.

   If you are copying a database using a full data backup only, you can either select the data backup from the backup catalog, or specify its location without using a backup catalog.

   A copy to a point in time is not possible if the full data backup is not recorded in a backup catalog.

Use a Backup Catalog	What Happens?
No	Copy the database without using a backup catalog. In the next steps, you are prompted to select the location of the data backup in the file system. The data backup to use is identified by its prefix (in a separate step).
Yes	Use a backup catalog to locate the data backup. In the next step, you are prompted to specify the location of the backup catalog.

4. Specify the backup location.

Backup Location	Subsequent Steps

File system	In the next step, specify the location and prefix of the data backup to be used.
	Example
	For the tenant database `PR1TENANT`, the location to specify could look like this:
	`/usr/sap/PR1/HDB00/backup/data/DB_PR1TENANT/`
	For the system database, the location to specify could look like this.
	`/usr/sap/PR1/HDB00/backup/data/SYSTEMDB/`

Backint	If you copy a database using a third-party backup tool, in the next step, specify the source system.

5. (Backint only) Specify the source database.

Source Database	Subsequent Steps
System database	Specify the SID of the source system.
Tenant database	Specify whether to copy the tenant database from: ○ A tenant database in the same SAP HANA system ○ A tenant database in a different SAP HANA system ○ A SAP HANA single-container system A SAP HANA single-container system can only be copied to a tenant database, not to a system database.

> **i** Note
>
> For a database copy, it is **not possible** to mix backups from different sources.
>
> The backup catalog, the data backups, and the log backups must be from either **only** a third-party tool or **only** the file system.
>
> (For a standard database recovery, it is possible to use a combination of backups from a third-party tool and the file system, provided that the backups originate from the same SAP HANA database.)

6. (File system only) Specify the backup to be used.

Backup Catalog	Backup to be Used
If you are using a backup catalog:	An overview of the available backups is displayed, depending on whether the catalog is in the file system or a third-party backup tool. It is not possible to mix backups from the different sources.
If you are not using a backup catalog:	In the next step, specify the location and the prefix of the backup. If required, specify alternative locations for the data and log backups.

7. (Optional) Check the availability of the backups to be used for the database copy.

   This check ensures that the backups exist at the location specified.
8. Review your settings.
   a. Choose *Review* to display a summary of the settings you specified.
   b. To display the SQL statement that will be used for the copy, choose *Display SQL Statement*.
   c. To change the settings, choose *Edit*. All the settings that you specified are retained until you change them.
9. Start the database copy.
   a. If the settings are correct, choose *Start Copy*.

      SAP HANA cockpit displays a warning that you are about to overwrite the target system.

b. To start the database copy, choose *Start Copy* again.

   The progress of the copy for each SAP HANA service is displayed.

   While the database is being copied, it is possible to cancel the copy process.

Results

When the database copy is completed, a message confirms this.

A copy of the SAP HANA database is created in the location you specified.

Database Copied	Next Steps
If you copied a system database	You now need to copy the tenant databases in the SAP HANA system.
If you copied a tenant database	The copy of the SAP HANA tenant database is now online and can be used by applications.

> **i** Note
>
> For a database copy to a point in time, SAP HANA cockpit shows the point in time to which the copy was made.
>
> The point in time that SAP HANA returns after a copy is the point in time of the last COMMIT to the database that has been copied.
>
> For this reason, this point in time may be before the point in time that you specified for the copy.

Database Credentials

To allow SAP HANA cockpit to connect to the copied database, you may need to change the credentials of the user that is registered in SAP HANA cockpit.

> **⚠** Caution
>
> Ensure that the correct passwords are used to connect to the copied SAP HANA database. If an incorrect password is used multiple times, SAP HANA may respond by locking that database user account.

Related Information

Working with Third-Party Backup Tools
Cancel a Recovery [page 635]
Recover SAP HANA From a Data Snapshot [page 632]

Important Disclaimer for Features in SAP HANA

For information about the capabilities available for your license and installation scenario, refer to the Feature Scope Description for SAP HANA.

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.
About the icons:

- Links with the icon  : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
  - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
  - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon : You are leaving the documentation for that particular SAP product or service and are entering a SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.
The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Gender-Related Language

We try not to use gender-specific word forms and formulations. As appropriate for context and readability, SAP may use masculine word forms to refer to all genders.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

www.sap.com/contactsap

© 2020 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see https://www.sap.com/about/legal/trademark.html for additional trademark information and notices.

THE BEST RUN SAP