PUBLIC

Document Version: 2024.8 – 2024-04-11

Administering SAP Datasphere

THE BEST RUN SAP

Content

1 Administering SAP Datasphere

Users with the *DW Administrator* role can configure, manage, and monitor the SAP Datasphere tenant to support the work of acquiring, preparing, and modeling data for analytics. They manage users and roles, create spaces, and allocate storage to them. They prepare and monitor connectivity for data integration and perform ongoing monitoring and maintainance of the tenant.

This topic contains the following sections:

> → Tip
>
> The English version of this guide is open for contributions and feedback using GitHub. This allows you to get in contact with responsible authors of SAP Help Portal pages and the development team to discuss documentation-related issues. To contribute to this guide, or to provide feedback, choose the corresponding option on SAP Help Portal:
>
> - *Feedback* ▶ *Edit page* ▶: Contribute to a documentation page. This option opens a pull request on GitHub.
> - *Feedback* ▶ *Create issue* ▶: Provide feedback about a documentation page. This option opens an issue on GitHub.
>
> You need a GitHub account to use these options.
>
> More information:
>
> - Contribution Guidelines
> - Introduction Video: Open Documentation Initiative
> - Blog Post: Introducing the Open Documentation Initiative

Configure Your SAP Datasphere Tenant

Either SAP will provision your tenant or you can create an instance in SAP BTP (see Creating and Configuring Your SAP Datasphere Tenant [page 16]).

- We recommend that you link your tenant to an SAP Analytics Cloud tenant (see Enable the Product Switch to Access an SAP Analytics Cloud Tenant [page 26]).
- You can enable SAP SQL data warehousing on your tenant to exchange data between your HDI containers and your SAP Datasphere spaces without the need for data movement (see Enable SAP SQL Data Warehousing on Your SAP Datasphere Tenant [page 27]).

- You can enable the SAP HANA Cloud script server to access the SAP HANA Automated Predictive Library (APL) and SAP HANA Predictive Analysis Library (PAL) machine learning libraries (see Enable the SAP HANA Cloud Script Server on Your SAP Datasphere Tenant [page 28]).

Create Users and Assign Roles

An administrator creates SAP Datasphere users manually, from a `*.csv` file, or via an identity provider (see Managing SAP Datasphere Users [page 45]).

You must assign one or more roles to each of your users via scoped roles and global roles (see Managing Roles and Privileges [page 52]). You can create your own custom roles or use the following standard roles delivered with SAP Datasphere:

- Roles providing privileges to administer the SAP Datasphere tenant:
  - **System Owner** - Includes all user privileges to allow unrestricted access to all areas of the application. Exactly one user must be assigned to this role.
  - **DW Administrator** - Can create users, roles and spaces and has other administration privileges across the SAP Datasphere tenant. Cannot access any of the apps (such as the *Data Builder*).
- Roles providing privileges to work in SAP Datasphere spaces:
  - **DW Space Administrator** (template) - Can manage all aspects of the spaces users are assigned to (except the *Space Storage* and *Workload Management* properties) and can create data access controls.
    - *DW Scoped Space Administrator* - This predefined scoped role is based on the DW Space Administrator role and inherits its privileges and permissions.

      > **Note**
      >
      > Users who are space administrators primarily need scoped permissions to work with spaces, but they also need some global permissions (such as Lifecycle when transporting content packages). To provide such users with the full set of permissions they need, they must be assigned to a scoped role (such as the *DW Scoped Space Administrator*) to receive the necessary scoped privileges, but they also need to be assigned directly to the *DW Space Administrator* role (or a custom role that is based on the *DW Space Administrator* role) in order to receive the additional global privileges.
  - *DW Integrator* (template) - Can integrate data via connections and can manage and monitor data integration in a space.
    - *DW Scoped Integrator* - This predefined scoped role is based on the DW Integrator role and inherits its privileges and permissions.
  - **DW Modeler** (template) - Can create and edit objects in the *Data Builder* and *Business Builder* and view data in objects.
    - *DW Scoped Modeler* - This predefined scoped role is based on the DW Modeler role and inherits its privileges and permissions.
  - **DW Viewer** (template) - Can view objects and view data output by views that are exposed for consumption in spaces.
    - *DW Scoped Viewer* - This predefined scoped role is based on the DW Viewer role and inherits its privileges and permissions.
- Roles providing privileges to consume the data exposed by SAP Datasphere spaces:

    - **DW Consumer** (template) - Can consume data exposed by SAP Datasphere spaces, using SAP Analytics Cloud, and other clients, tools, and apps. Users with this role cannot log into SAP Datasphere. It is intended for business analysts and other users who use SAP Datasphere data to drive their visualizations, but who have no need to access the modeling environment.
        - *DW Scoped Consumer* - This predefined scoped role is based on the DW Consumer role and inherits its privileges and permissions.
- Roles providing privileges to work in the SAP Datasphere catalog:
    - **Catalog Administrator** - Can set up and implement data governance using the catalog. This includes connecting the catalog to source systems for extracting metadata, building business glossaries, creating tags for classification, and publishing enriched catalog assets so all catalog users can find and use them. Must be used in combination with another role such as *DW Viewer* or *DW Modeler* for the user to have access to SAP Datasphere.
    - **Catalog User** - Can search and discover data and analytics content in the catalog for consumption. These users may be modelers who want to build additional content based on official, governed assets in the catalog, or viewers who just want to view these assets. Must be used in combination with another role such as *DW Viewer* or *DW Modeler* for the user to have access to SAP Datasphere.

Create Spaces and Allocate Storage to Them

All data acquisition, preparation, and modeling in SAP Datasphere happens inside spaces. A space is a secure area - space data cannot be accessed outside the space unless it is shared to another space or exposed for consumption.

An administrator must create one or more spaces. They allocate disk and memory storage to the space, set its priority, and can limit how much memory and how many threads its statements can consume. See Creating Spaces and Allocating Storage [page 99].

Prepare Connectivity

Administrators prepare SAP Datasphere for creating connections to source systems in spaces (see Preparing Connectivity for Connections [page 109]).

Monitor and Maintain SAP Datasphere

Administrators have access to various monitoring logs and views and can, if necessary, create database analysis users to help troubleshoot issues (see Monitoring SAP Datasphere [page 188]).

1.1 Administration Apps and Tools

You administer SAP Datasphere using apps and tools in the side navigation area.

(*Space Management*)

In the *Space Management*, you can set up, configure, and monitor your spaces, including assigning users to them. For more information, see Preparing Your Space and Integrating Data.

(*System Monitor*)

In the *System Monitor*, you can monitor the performance of your system and identify storage, task, out-of-memory, and other issues. For more information, see Monitoring SAP Datasphere [page 188].

Security

Tool	Task	More Information
Users	Create, modify, and manage users in SAP Datasphere.	Managing SAP Datasphere Users [page 45]
Roles	Assign pre-defined standard roles or custom roles that you have created to users.	Managing Roles and Privileges [page 52]
Activities	Track the activities that users perform on objects such as spaces, tables, views, data flows, and others, track changes to users and roles, and more.	Monitor Object Changes with Activities [page 207]

System > Configuration

Tab	Task	More Information
Data Integration	*Live Data Connections (Tunnel)*: For SAP BW/4HANA model import, you need Cloud Connector to make http requests to SAP BW/4HANA. This requires a live data connection of type tunnel to SAP BW/4HANA.	Create Live Data Connection of Type Tunnel [page 151]
	On-Premise Agents: Manage Data Provisioning Agents which are required to act as gateway to SAP Datasphereto enable using connections to on-premise sources for remote tables and building views.	Connect and Configure the Data Provisioning Agent [page 117] Register Adapters with SAP Datasphere [page 119] Monitoring Data Provisioning Agent in SAP Datasphere [page 163] Pause Real-Time Replication for an Agent [page 167]
	Third-Party Drivers: Upload driver files that are required for certain third-party cloud connections to use them for data flows.	Upload Third-Party ODBC Drivers (Required for Data Flows) [page 132]
Tenant Links	*Link My Tenants*: Link an SAP Analytics Cloud tenant to your SAP Datasphere tenant to enable the product switch in the top right of the shell bar, and be able to easily navigate between them.	Enable the Product Switch to Access an SAP Analytics Cloud Tenant [page 26]
Security	*SSL/TLS Certificates* : Upload server certificates to enable secure SSL/TLS-based connections to certain sources.	Manage Certificates for Connections [page 131]
	Password Policy Configuration: Define your password policy settings for the database users. The policy can be enabled when configuring your database users.	Set a Password Policy for Database Users [page 52]
Audit	*Audit View Enablement*: Configure a space that gets access to audit views and allows you to display the audit logs in that space.	Enable Audit Logging
	Audit Log Deletion:	Delete Audit Logs [page 206]
Monitoring	*Monitoring View Enablement*:	Analyze Monitoring Data in a Space [page 197]
	Expensive Statement Tracing:	
	MDS Information Tracing:	

Tab	Task	More Information
IP Allowlist	*Trusted IPs*: Control the range of external public IPv4 addresses that get access to the database of your SAP Datasphere by adding them to an allowlist.	Add IP address to IP Allowlist [page 128]
	Trusted Cloud Connector IPs:	
Tasks	Clean-up task logs to reduce storage consumption in your SAP Datasphere tenant. Also allows you to view a list of users whose authorization consent will expire within a given timeframe, by default, four weeks.	Deleting Task Logs to Reduce Storage Consumption Check Consent Expirations [page 211]
Database Access	*Database Analysis Users*: Create a database analysis user to connect to your SAP HANA Cloud database to analyze, diagnose and solve database issues. Only create this user for a specific specific task and delete right after the task has been completed.	Monitoring SAP Datasphere [page 188]
	Database User Groups: Create an isolated environment with corresponding administrators where you can work more freely with SQL in your SAP HANA Cloud database.	Creating a Database User Group [page 180]
Tenant Configuration	Allocate the capacity units to storage and compute resources for your tenant.	Configure the Size of Your SAP Datasphere Tenant [page 21]
SAP BW Bridge		Provisioning the SAP BW Bridge Tenant

System Administration

Tab	Task	More Information
System Configuration	*Session timeout*: Set the amount of time before a user session expires if the user doesn't interact with the system.	By default the session timeout is set to 3600 seconds (1 hour). The minimum value is 300 seconds, and the maximum value is 43200 seconds.
	Allow SAP support user creation: Let SAP create support users based on incidents. Support users generated by SAP will be deleted after their validity has expired or after the incident has been closed.	Request Help from SAP Support [page 14]

Tab	Task	More Information
Data Source Configuration	*SAP Cloud Platform (SAP CP) Account*: Get subaccount information for SAP Datasphere. You need the information to configure the Cloud Connector that SAP Datasphere uses to connect to sources for data flows and model import.	Set Up Cloud Connector in SAP Datasphere [page 127]
	Live Data Sources: If you want to use SAP BW/4HANA model import, you need to allow data from your live data connection of type tunnel to securely leave your network.	
	On-premise data sources: Add location IDs if you have connected multiple Cloud Connector instances to your SAP Datasphere subaccount and you want to offer them for selection when creating connections using a Cloud Connector.	
Security	*Authentication Method*: Select the authentication method used by SAP Datasphere.	Enabling a Custom SAML Identity Provider [page 36]
	SAML Single Sign-On (SSO) Configuration: Configure SAML SSO if you selected it as authentication method.	
App Integration	*OAuth Clients*: You can use Open Authorization (OAuth) protocol to allow third-party applications access.	Create OAuth2.0 Clients to Authenticate Against SAP Datasphere [page 29]
	Trusted Identity Providers: If you use the OAuth 2.0 SAML Bearer Assertion workflow, you must add a trusted identity provider.	
	Trusted Origins: Enter the origins that will be hosting your client application.	
Notifications	Make sure that users are notified appropriately about issues in the tenant.	Configure Notifications [page 210]

System >> About

Every user can view information about the software components and versions of your system, in particular:

- *Version*: Displays the version of the SAP Datasphere tenant.
- *Build Date*: Displays the date and time when the current version of the SAP Datasphere tenant was built.
- *Tenant*: Displays the SAP Datasphere tenant id.
- *Database*: Displays the id of the SAP Datasphere run-time database.

- *Platform Version*: Displays the version of the SAP Analytics Cloud components used in SAP Datasphere.

Users with the DW Administrator role can open a *More* section to find more details. They can find outbound and database IP addresses that might be required for allowlists in source systems or databases of SAP Datasphere for example (see Finding SAP Datasphere IP addresses [page 130]).

1.2 System Requirements and Technical Prerequisites

SAP Datasphere is a fully web-based offering. You will need an internet connection and a system that meets certain requirements.

The requirements listed here are for the current release.

Client Software Requirements

Client Software	Version	Additional Information
Desktop browser	Google Chrome, latest version	Google releases continuous updates to their Chrome browser. We make every effort to fully test and support the latest versions as they are released. However, if defects are introduced with OEM-specific browser software, we cannot guarantee fixes in all cases. For additional system requirements, see your web browser documentation.
	Microsoft Edge based on the Chromium engine, latest version	Microsoft has available for download continuous updates to their new Chromium-based Edge browser. We make every effort to fully test and support the latest versions as they are released.
Additional software	Adobe Acrobat Reader 9 or higher	-

Client Configuration Requirements

Client Configuration	Setting	Additional Information
Network bandwidth	Minimum 500-800 kbit/s per user	In general, SAP Datasphere requires no more bandwidth than is required to browse the internet. All application modules are designed for speed and responsiveness with minimal use of large graphic files.

Client Configuration	Setting	Additional Information
Screen resolution	XGA 1024x768 (high color) or higher Widescreen: 1366x766 or higher	-
Minimum recommended browser cache size	250 MB	SAP Datasphere is a Web 2.0 application. We recommend allowing browser caching because the application uses it heavily for static content such as image files. If you clear your cache, the browser will not perform as well until the deleted files are downloaded again to the browser and cached for use next time. To set browser cache size, see your browser documentation.
HTTP 1.1	Enable	-
JavaScript	Enable	-
Cookies	Enable web browser session cookies (non-persistent) for authentication purposes	-
Pop-up windows	Allow pop-up windows from SAP Datasphere domains	-
Power Option Recommendation	High Performance mode for improved JavaScript performance	For Microsoft based Operating Systems

Supported Languages

Client Browser	What's Supported
Menus, buttons, messages, and other elements of the user interface.	Bulgarian (bgBG); Catalan (caES); Chinese (zhTW); Chinese (Simplified) (zhCN); Croatian (hrHR); Czech (csCZ); Danish (daDK); Dutch (nlNL); English (enGB); English (enUS); Estonian (etEE); French (frCA); French (frFR); Finnish (fiFI); German (deDE); German (deCH); Greek (elGR); Hindi (hiIN); Hungarian (huHU); Indonesian (idID); Italian (itIT); Japanese (jaJP); Korean (koKR); Latvian (lvLV); Lithuanian (ltLT); Malay (msMY); Norwegian (noNO); Polish (plPL); Portuguese (Brazil) (ptBR); Portuguese (Portugal) (ptPT); Romanian (roRO); Russian (ruRU); Serbian (srRS); Slovakian (skSK); Slovenian (slSL); Spanish (esES); Spanish (esMX); Swedish (svSE); Thai (thTH); Turkish (trTR);Ukrainian (ukUA); Vietnamese (viVN) and Welsh (cyGB).

Data Connectivity

Connectivity with SAP HANA Smart Data Integration

We recommend to always use the latest released version of the Data Provisioning Agent but at least the recommended minimum version from SAP Note 2419138. Make sure that all agents that you want to connect to SAP Datasphere have the same latest version.

For more information, including information on minimum requirements for source systems and databases, see:

- SAP HANA Smart Data Integration Product Availability Matrix (PAM)
- Configure Data Provisioning Adapters in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation

1.3 Request Help from SAP Support

You can request help from SAP Product Support by creating a support incident. In many cases, a support user is required to allow an SAP support engineer to log into and troubleshoot your system.

You can create an SAP support incident on the SAP Support Portal (S-user login required). For detailed information about what to include in an incident, see SAP Note 2854764.

An administrator can make sure that a support user is created in your tenant. Two options are available:

- An administrator creates the support user.
  Before creating an incident with SAP, proceed as follows:
  1. In the shell bar, click (*Support*).
  2. In the *Support* dialog, click *Create Support User* and then choose *OK* to confirm the support user creation.
     An email is automatically sent to SAP Support to notify them of the newly created support user, and it is listed with your other users at *Security* > *Users*.
     The support user has minimum privileges and does not consume a user license.
     You can assign an appropriate role to the support user (DW Administrator role) and add it to the required space.
  3. Delete the support user when your issue is resolved.
- An administrator generally allows SAP Product Support to create support users based on incidents.
  Proceed as follows:
  1. In the side navigation area, click (*System*) > (*Administration*) > *System Configuration*.

     > **Note**
     >
     > If your tenant was provisioned prior to version 2021.03, click *(Product Switch)* > *Analytics* > *System* > *Administration* > *System Configuration*.

  2. Choose *Edit*.
  3. Set the *Allow SAP support user creation* setting to *ON*.
  4. Click *Save*.
     In case of an incident, a support engineer from SAP Product Support can now request and generate a personalized support user for the affected tenant. This user is enabled for multi-factor authentication.

The support user either has read-only privileges (DW Viewer role) or, if requested by the support engineer and confirmed by the customer in the incident, the support user can have the DW Administrator role. The user does not consume a user license.
The support user will be automatically deleted after two days or after the incident has been closed.

For more information about creating a support user, see SAP Note 2891554.

2 Creating and Configuring Your SAP Datasphere Tenant

You can create your own tenant in the SAP BTP Cockpit. The procedure is the same for both subscription-based and consumption-based contracts. Some details may vary depending on the chosen service plan (free or standard). For more information about limitations for a free plan, see SAP Note 3227267.

Characteristics	Standard Plan	Free Plan
Provisioning	• For information about region availability, see the SAP Discovery Center. • The SAP BTP subaccount administrator must trigger the SAP Datasphere instance creation. Tenant creation will not be triggered by SAP. • You must create and configure the to-be-provisioned SAP Datasphere service instance (tenant) in SAP BTP. See Create Your SAP Datasphere Service Instance in SAP BTP [page 19]. • The system owner of SAP Datasphere, who has been specified during the provisioning, is notified via email when the tenant is provisioned.	• For information about region availability, see the SAP Discovery Center. • The SAP BTP subaccount administrator must trigger the SAP Datasphere instance creation. Tenant creation will not be triggered by SAP. • You must create and configure the to-be-provisioned SAP Datasphere service instance (tenant) in SAP BTP. See Create Your SAP Datasphere Service Instance in SAP BTP [page 19]. • The system owner of SAP Datasphere, who has been specified during the provisioning, is notified via email when the tenant is provisioned.
Size Configuration **Note** For maximum size configuration options, see the tables below.	• Tenants are initially created with minimal configuration that includes 256 GB of storage and 2 compute blocks. • Once logged to your tenant, upscaling can be done at any time. See Configure the Size of Your SAP Datasphere Tenant [page 21]. **Note** After finalizing the configuration, you can only change the size of your SAP BW Bridge storage later if you don't have any SAP BW Bridge instances. To view all supported size combinations, go to the SAP Datasphere Capacity Unit Estimator.	• The minimal and only possible configuration provided for SAP Datasphere resources is: 128 GB of storage and 1 compute block. • You cannot upscale free plan tenants. You need to update your plan from free to standard if any sizing configuration is required.
Metering	• The number of consumed capacity units is reported on a hourly basis to your SAP BTP account.	• The usage of a free plan tenant is reported to your SAP BTP account, but SAP does not charge you for using this tenant.

Characteristics	Standard Plan	Free Plan
Time Limitation	• **Subscription Contract**: dependent on the contract. • **Consumption Based Contract**: no time limitation.	: The time limitation is 90 days and the trial duration cannot be extended. You can update the tenant from free to standard plan before the 90-day expiration (the number of days before the expiration is displayed in the top panel of the SAP Datasphere free-plan tenant). If you do not perform the update within 90 days, the tenant is automatically deleted. The remaining service instance cannot be reused and should be deleted at any time by an administrator of your SAP BTP account. See Update Your Free Plan to Standard Plan in SAP BTP [page 25].
Number of tenants	No limitation. In the case of a subscription contract, the available capacity units can be distributed among all your tenants.	1 per SAP BTP global account.

Maximum Configuration Values

The maxium configuration size of your tenant depends on regional availability and your server type.

> **Note**
>
> - Data integration includes 200h/month from the minimum free package.
> - Catalog includes 0.5 GB/h from the minimum free package.
> - Catalog storage/crawling includes 100h/month from the minimum free package.

Amazon Web Services (AWS)

Hyperscaler Regional Availability	Compute	Storage	BW Bridge	Data Lake	Data Integration	Catalog	Catalog Storage / Crawling
Australia	1800 GB	7168 GB	4096 GB	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month
Brazil (São Paulo)	1800 GB	7168 GB	4096 GB	Not Supported	7200 h/ month	20.5 GB/h	2100 h/ month
Canada (Montreal)	1800 GB	7168 GB	4096 GB	Not Supported	7200 h/ month	20.5 GB/h	2100 h/ month

Hyperscaler Regional Availability	Compute	Storage	BW Bridge	Data Lake	Data Integration	Catalog	Catalog Storage / Crawling
Europe (Frankfurt)	5970 GB	15872 GB	4096 GB	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month
EU Access (Frankfurt)	5970 GB	15872 GB	4096 GB	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month
Japan (Tokyo)	1800 GB	7168 GB	4096 GB	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month
Singapore	1800 GB	7168 GB	4096 GB	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month
South Korea	1800 GB	7168 GB	4096 GB	Not Supported	7200 h/ month	20.5 GB/h	2100 h/ month
US East	5970 GB	15872 GB	4096 GB	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month

Microsoft Azure

Hyperscaler Regional Availability	Compute	Storage	BW Bridge	Data Lake	Data Integration	Catalog	Catalog Storage / Crawling
Europe (Amsterdam)	5600 GB	13824 GB	Supported	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month
Europe (Switzerland)	5600 GB	13824 GB	Supported	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month
US West	5600 GB	13824 GB	Supported	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month

Google Cloud Platform (GCP)

Hyperscaler Regional Availability	Compute	Storage	BW Bridge	Data Lake	Data Integration	Catalog	Catalog Storage / Crawling
Europe (Frankfurt)	1344 GB	3328 GB	Supported	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month
India (Mumbai)	1344 GB	3328 GB	Supported	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month
US Central	5750 GB	14336 GB	Supported	90 TB	7200 h/ month	20.5 GB/h	2100 h/ month

2.1 Create Your SAP Datasphere Service Instance in SAP BTP

Create your SAP Datasphere service instance in SAP Business Technology Platform.

> **Note**
>
> Creating an SAP Datasphere service instance in SAP Business Technology Platform (SAP BTP) results in provisioning an SAP Datasphere tenant.

For both subscription-based contracts (initiated on November 2023) and consumption-based contracts, you can access the SAP BTP cockpit and view all currently available services in a global account. You need to structure this global account into subaccounts and other related artefacts, such as directories and/or spaces.

Prerequisites

To create your SAP Datasphere service instance in SAP BTP, you need the following prerequisites:

- Your global account has a commercial entitlement either via cloud credits (in case of a consumption-based contract) or via a subscription-based contract.
- A Cloud Foundry subaccount which is entitled for SAP Datasphere. For more information, see Configure Entitlements and Quotas for Subaccounts.
- You have SAP BTP administration authorization on the subaccount that is entitled to SAP Datasphere.
- You are using Google Chrome to properly view popups in SAP BTP.

Service Plans

Service Plan	Description
Standard	The standard plan provides an SAP Datasphere tenant for productive and non-productive use, which is represented by a service instance.
Free	The free plan provides an SAP Datasphere tenant for a limited time for trial use, which is represented by a service instance.

> **Note**
>
> For information about region availability, see the SAP Discovery Center.

Create a Tenant

The following procedure uses the SAP BTP cockpit to create the service instance.

In the *SAP Datasphere Administration Guide*, we provide high-level steps to create an SAP Datasphere tenant on SAP BTP. For more detailed information, or for instructions that use the Cloud Foundry Command-Line Interface, see the SAP Business Technology (SAP BTP) documentation.

> **Note**
>
> You can create only one free tenant under the global account. If your SAP BTP service causes issues, you can open an incident ticket via ServiceNow.

1. In the SAP BTP cockpit, navigate to the space in which you want to create the service instance, and click *Services* > *Service Marketplace* in the left navigation area.
   For more information, see Navigate to Orgs and Spaces.
2. Search for "Datasphere", and click the SAP Datasphere service to open it.
3. Click *Create* in the top-right corner.
   A wizard opens, in which you can select or specify the following parameters:

Parameter	Description
Service	Select SAP Datasphere.
Plan	Select *Standard* or *Free*.
Runtime Environment	Select a runtime environment. **Note** Not all runtime environments are available for free.
Space	[no selection needed if you're creating the instance from the space area] Select the SAP BTP space in which you want to create the service instance.
Instance Name	Enter a name to identify your instance (up to 32 alphanumeric characters, periods, underscores, and hyphens; cannot contain white spaces).

4. Click *Next* and enter the following information about the SAP Datasphere system owner, who will be notified when the service instance is created: *First Name*, *Last Name*, *Email*, and *Host Name*.

   > **Note**
   >
   > Alternatively, you can use a JSON file to provide the information above.

5. Click *Next* to go to the final page of the wizard where you can review your selections, and then click *Create* to exit the wizard.
   An information message is displayed to confirm that the service instance creation is in progress.

   > **Note**
   >
   > The creation of the instance can take a while.

6. Click *View Instance* to go to your space *Service Instances* page, where the new instance is listed and you can view the progress of its creation.
7. When the service instance is created, the SAP Datasphere system owner receives an email confirming its availability, and providing a link to navigate to the SAP Datasphere tenant, which the service instance represents.

> ⓘ Note
>
> If the creation of the service instance fails (the "failed" status is displayed), you must first delete the failed instance and then create a new SAP Datasphere service instance. If you need support, you can open an incident via ServiceNow with the component DS-PROV.

2.2 Configure the Size of Your SAP Datasphere Tenant

Configure the size of your tenant by specifying resource sizes based on your business needs. Capacity Units (CU) are allocated to obtain storage and compute resources for your tenant.

You can configure the size of a subscription-based tenant and a consumption-based tenant with a standard plan.

To do so, you must have an SAP Datasphere administrator role.

Configuring the Sizes of Resources

In the *Tenant Configuration* page of the *Configuration* area, you can increase the sizes for the various resources, within the permitted size combinations, to obtain a configuration that fits your exact needs, and then click *Save*.

> ⚠ Caution
>
> Once you save the size configuration of your tenant, be aware that some resources cannot be resized later. Storage cannot be downsized. If you require a storage downsize, you must recreate the tenant. Exception: If you need to decrease the memory, see SAP note 3224686.

Also, once you click *Save*:

- The whole process may take more than 90 minutes. The configuration process is not long, but the operational process in the background can take a while.
- In case an error occurs, you are notified that the configuration cannot be completed and that you need to try again later by clicking the *Retry* button (which replaces the *Save* button in such a case). The delay depends on the error (for example, if there is an error on the SAP HANA Cloud database side, you may need to retry after 60 minutes).
- You can only make changes to SAP HANA Compute and SAP HANA Storage once every 24 hours.
- If you try to change your SAP HANA configuration, SAP HANA Cloud functionalities (Spaces, DPServer, Serving of Queries) will not be available for around 10 minutes. If you run into issues after the configuration, use the *Retry* button.

Supported Sizes

To view all supported size combinations for compute and storage resources and the number of capacity units consumed, go to the SAP Datasphere Capacity Unit Estimator.

Tenant Configuration Page Properties

Base Configuration

Property	Description
Performance Class	Select a performance class for your tenant: • Memory • Compute • High-Memory • High-Compute **Note** The performance class you select determines the number of vCPUs allocated to your tenant.
Storage	Set the size of disk storage. You can specify from 128 GB (minimum), by increments of 64 GB.
Memory	Set the memory allocated to your tenant. You can specify from 128 GB (minimum), by increments of 64 GB.
vCPU	Displays the number of vCPUs allocated to your tenant. The number is calculated based on the selected performance class, and memory used by your tenant.
Enable the SAP HANA Cloud Script Server	Enable this to access the SAP HANA Automated Predictive Library (APL) and SAP HANA Predictive Analysis Library (PAL) machine learning libraries.

Additional Data Warehouse Configuration

Property	Description
Data Lake Storage	[optional] Select the size of data lake disk storage. You can specify from 0 TB (minimum) to 90 TB (maximum), by increments of 1 TB. Data lake storage includes data lake compute. To reduce the size of your data lake storage, you must first delete your data lake instance, and re-create it in the size that you want. **Note** Deletion cannot be reversed and all data stored in the data lake instance will be deleted. You cannot delete your data lake storage if it's connected to a space. You must first disconnect the space: 1. Go to *Space Management*, choose a space. 2. Select *Edit*. 3. Under *General Settings*, clear the *Use this space to access data lake* checkbox. Data lake is not available in all regions. See SAP note 3144215.
SAP BW Bridge Storage	[optional] Select the size of SAP BW bridge using the dropdown menu, starting from 0 GB (minimum). SAP BW Bridge includes SAP BTP, ABAP environment, runtime and compute. **Caution** You can only change the SAP BW Bridge storage allocation, if no SAP BW Bridge instances are created. The process for allocating capacity units to SAP BW Bridge is not part of the configuration process. **Note** • When you provision a non-productive tenant, you'll need to request the SAP BW bridge Cloud-ABAP Service for your tenant by opening an incident via ServiceNow with the component DWC-BWB. • First finalize the size configuration of your tenant, then open the incident as a next step. Once the incident has been processed, you can create the SAP BW bridge instance in the dedicated page *SAP BW Bridge* of the *Configuration* area with the size you've allocated (see Provisioning the SAP BW Bridge Tenant). • SAP BW Bridge is not available in all regions. See SAP note 3144215. • When using a test tenant with the minimal configuration (128 GB of storage and 1 compute block), you cannot add SAP BW Bridge storage unless you add more compute blocks. • As soon as you click *Save*, the allocated capacity units will be assigned to SAP BW Bridge.

Data Integration

Property	Description
Data Integration	[optional] Select the number of compute blocks to allocate to Data Integration applications which provide enhanced data integration capabilities. This comprises the replication flow that enables data integration with delta processing which is the recommended integration solution for all supported source and target systems. You can increase the number of blocks to assign a maximum of 7200 node hours. You can decrease the number of blocks until you reach the minimum number of node hours included in your plan. **Note** Data flows are not part of this commercial package.
Execution Hours	The number of node hours available for Data Integration applications per month is calculated from the number of allocated compute blocks. Every month you're entitled to running up to 200 hours of jobs. Using more than 200 hours has no impact in other jobs. The consumption of data integration is not limited to avoid interrupting critical integration scenarios, but you will be billed for the overusage when it happens.
Maximum Parallel Jobs	The maximum number of parallel jobs is calculated from the number of execution hours assigned. For every 100 execution hours, you are given one extra parallel job, up to a maximum of 10. Each parallel job means that roughly 5 datasets, from one or several replication flows, can be run in parallel. If more replication flows are running, processing will be queued and replication will occur less frequently.
Data Integration: *Allocated Execution Hours*	Displays the number of hours allocated to Data Integration applications.
Data Integration: *Used Execution Hours*	Displays the number of hours used by Data Integration applications.
Data Integration: *Exceeded Execution Hours*	Displays the execution hours you have used that exceed the amount allocated by your tenant configuration. **Note** This only appears if you have used more hours than allocated.

Premium Outbound Integration

Property	Description
Outbound Blocks	You can increase or decrease the number of storage blocks allocated for premium outbound integration. Each block provides 20 GB of storage.
Outbound Volume	The outbound volume is calculated from the number of allocated blocks.
Premium Outbound Usage: *Allocated Data Volume*	Displays the number of GB allocated to premium outbound integration.

Property	Description
Premium Outbound Usage: *Used Data Volume*	Displays the number of GB used by premium outbound integration.
Premium Outbound Usage: *Exceeded Data Volume*	Displays the data volume you have used that exceeds the amount allocated by your tenant configuration. **Note** This only appears if you have used more data than allocated.

Catalog

Property	Description
Catalog Storage	Included by default. You can increase or decrease the number of storage blocks allocated for the catalog.
Storage	The amount of storage available for the catalog is calculated from the number of allocated blocks.
Catalog Usage: *Allocated Storage*	Displays the number of GB allocated to the catalog.
Catalog Usage: *Used Storage*	Displays the number of GB used by the catalog.
Catalog Usage: *Exceeded Storage*	Displays the amount of storage you have used that exceeds the amount allocated by your tenant configuration. **Note** This only appears if you have used more storage space than allocated.

Capacity Units

Property	Description
Units in Use per Month	Displays the estimated number of capacity units consumed per month by the storage and compute resources you've specified.
Units in Use per Hour	Displays the estimated number of capacity units consumed per hour by the storage and compute resources you've specified.

2.3 Update Your Free Plan to Standard Plan in SAP BTP

Update your service instance from free plan to standard plan.

In SAP Business Technology Platform (SAP BTP), if you have an SAP Datasphere service instance with a free plan, which you can use for 90 days, you can update it to a standard plan (no time limitation) for productive purposes. The number of days before the expiration is displayed in the top panel of SAP Datasphere.

Note

If you do not update to a standard plan within 90 days, your SAP Datasphere tenant will be suspended. While the tenant is suspended, you can still upgrade your service instance from the free to standard plan, but after 5 days of suspension, your tenant will be deleted and there is no way to recover it.

If your tenant is deleted, the service instance will still be shown in your Global Account, but it is not functional. You can delete it and create a new SAP Datasphere service instance with a free plan.

To do so, you must have SAP BTP administration authorization on the subaccount that is entitled to SAP Datasphere.

1. In SAP BTP, select the subaccount and the space where the service instance with a free plan was created.
2. Navigate to *Instances and Subscriptions*.
3. In the *Service Instances* page, find the SAP Datasphere service instance with the free plan, click the button at the end of the row and select *Update*.

   > **Note**
   >
   > After updating your free plan to standard plan, you must wait at least 24 hours before changing the tenant settings on the *Tenant Configuration* page.

4. In the *Update Instance* dialog, select *standard* and click *Update Instance*.
   You can view the progress of the update. The status of the instance becomes green when the update is completed.

   > **Note**
   >
   > The update process takes around 30 minutes, and during this time some features might not work as expected.

2.4 Enable the Product Switch to Access an SAP Analytics Cloud Tenant

You can link your SAP Datasphere tenant to a SAP Analytics Cloud tenant and enable the product switch in the top right of the shell bar, to help your users easily navigate between them.

Procedure

1. Go to *System* > *Configuration* > *Tenant Links*.
2. Enter the URL of your SAP Analytics Cloud tenant.
3. Click *Save* to confirm the connection.

   > **Note**
   >
   > An SAP Analytics Cloud user must create a live connection before they can consume data from SAP Datasphere (see Live Data Connections to SAP Datasphere in the *SAP Analytics Cloud* documentation).
   >
   > Multiple SAP Analytics Cloud tenants can create live connections to your SAP Datasphere tenant, but only one SAP Analytics Cloud tenant can be accessed via the product switch.

For more information about consuming data in this way, see Consume Data in SAP Analytics Cloud via a Live Connection.

2.5 Enable SAP SQL Data Warehousing on Your SAP Datasphere Tenant

Use SAP SQL Data Warehousing to build calculation views and other SAP HANA Cloud HDI objects directly in your SAP Datasphere run-time database and then exchange data between your HDI containers and your SAP Datasphere spaces. SAP SQL Data Warehousing can be used to bring existing HDI objects into your SAP Datasphere environment, and to allow users familiar with the HDI tools to leverage advanced SAP HANA Cloud features.

Context

To enable SAP SQL Data Warehousing on your SAP Datasphere tenant, an S-user must create an SAP ticket to connect your SAP BTP account.

> **Note**
>
> The SAP Datasphere tenant and SAP Business Technology Platform organization and space must be in the same data centre (for example, eu10, us10).

For information about using this feature, see Exchanging Data with SAP SQL Data Warehousing HDI Containers.

Procedure

1. In the side navigation area, click  (*Space Management*), locate your space tile, and click *Edit* to open it.
2. In the *HDI Containers* section, click *Enable Access* and then click *Open Ticket* to create an SAP ticket for the `DWC-SM` component to request us to map your SAP Datasphere tenant to your SAP Business Technology Platform account.
3. Provide the following information:

Item	Description
SAP Datasphere Tenant ID	In the side navigation area, click (*System*) > (*About*). **Note** You need the *Tenant ID* for the ticket, and the *Database ID* when building your containers in the SAP Datasphere run-time database.
SAP Business Technology Platform Org GUID	Your SAP Business Technology Platform organization ID. You can use the Cloud Foundry CLI to find your organization GUID (see https://cli.cloudfoundry.org/en-US/v6/org.html).
SAP Business Technology Platform Space GUID	The SAP Business Technology Platform space inside the organization. You can use the Cloud Foundry CLI to find your space GUID (see https://cli.cloudfoundry.org/en-US/v6/space.html).

You will be notified when your ticket has been processed.

4. Build one or more new HDI containers in the SAP Datasphere run-time database (identified by the *Database ID* on the SAP Datasphere *About* dialog).

   For information about setting up your build, see Set Up an HDI Container .
5. When one or more containers are available in the run-time database, the *Enable Access* button is replaced by the + button in the *HDI Containers* section for all your SAP Datasphere spaces (see Add an HDI Container and Access its Objects in Your Space).

2.6 Enable the SAP HANA Cloud Script Server on Your SAP Datasphere Tenant

You can enable the SAP HANA Cloud script server on your SAP Datasphere tenant to access the SAP HANA Automated Predictive Library (APL) and SAP HANA Predictive Analysis Library (PAL) machine learning libraries.

To enable the SAP HANA Cloud script server, go to the *Tenant Configuration* page and select the checkbox in the *Base Configuration* section. For more information, see Configure the Size of Your SAP Datasphere Tenant [page 21].

Note

The script server cannot be enabled in a SAP Datasphere consumption-based tenant with free plan.

Once the script server is enabled, the *Enable Automated Predictive Library and Predictive Analysis Library* option can be selected when creating a database user (see Create a Database User).

For detailed information about using the machine learning libraries, see:

- SAP HANA Automated Predictive Library Developer Guide

- SAP HANA Cloud Predictive Analysis Library (PAL)

2.7 Create OAuth2.0 Clients to Authenticate Against SAP Datasphere

Users with the *DW Administrator* role can create OAuth2.0 clients and provide the client parameters to users who need to connect clients, tools, or apps to SAP Datasphere.

Context

Users that have access to an OAuth client can:

- Log into the Command Line Interface via an OAuth Client
- Consume SAP Datasphere Data in SAP Analytics Cloud via an OData Service
- Consume Data in Power BI and Other Clients, Tools, and Apps via an OData Service

> **Note**
>
> Consuming exposed data in third-party clients, tools, and apps via an OData service requires a three-legged OAuth2.0 flow with type `authorization_code`. For information about creating an OAuth client, see Create OAuth2.0 Clients to Authenticate Against SAP Datasphere [page 29].

Procedure

1. In the side navigation area, click (*System*) (*Administration*) *App Integration*.
2. Under *Configured Clients*, select *Add a New OAuth Client*.
3. In the dialog, enter or review the following properties as appropriate:

Property	Description
Name	Enter a name to identify the OAuth client.
OAuth Client ID	[read-only] Displays the ID once the client is created.
Purpose	Select *Interactive Usage*. **Note** The *API Access* purpose is not used in SAP Datasphere.
Authorization Grant	[read-only] Only the *Authorization Code* grant is available.

Property	Description
Secret	[read-only] Allows the secret to be copied immediately after the client is created. **Note** Once you close the dialog, the secret is no longer available.
Redirect URI	Enter a URI to indicate to where the user will be redirected after authorization. If the URI has dynamic parameters, use a wildcard pattern (for example, `https://redirect_host/**`). The client, tool, or app that you want to connect is responsible for providing the redirect URI. When creating an OAuth client to allow SAP Analytics Cloud to connect to SAP Datasphere via an OData services connection (see Consume SAP Datasphere Data in SAP Analytics Cloud via an OData Service), use the *Redirect URI* provided in the SAP Analytics Cloud connection dialog.
Token Lifetime	Enter a lifetime for the access token from a minimum of 60 seconds to a maximum of one day. Default: 60 minutes
Refresh Token Lifetime	Enter a lifetime for the refresh token from a minimum of 60 seconds to a maximum of 180 days. Default: 30 days

4. Click *Add* to create the client and generate the ID and secret.
5. Copy the secret, save it securely, and then close the dialog.

> **Note**
> You won't be able to copy the secret again. If you lose it, you will need to create a new client.

6. Provide the following information to users who will use the client:

Standard OAuth2 Authorization Flow	OAuth2SAMLBearer Principal Propagation Flow
• *Client ID* • *Secret* • *Authorization URL* • *Token URL*. Users must manually authenticate against the IDP in order to generate the authorization code before continuing with the remaining OAuth2.0 steps.	• *Client ID* • *Secret* • *OAuth2SAML Token URL* • *OAuth2SAML Audience* Users authenticate with their third-party app, which has a trusted relationship with the IDP, and do not need to re-authenticate (see Add a Trusted Identity Provider [page 31]).

2.7.1 Add a Trusted Identity Provider

If you use the OAuth 2.0 SAML Bearer Assertion workflow, you must add a trusted identity provider to SAP Datasphere.

Context

The OAuth 2.0 SAML Bearer Assertion workflow allows third-party applications to access protected resources without prompting users to log into SAP Datasphere when there is an existing SAML assertion from the third-party application identity provider.

> **Note**
>
> Both SAP Datasphere and the third-party application must be configured with the same identity provider. The identity provider must have a user attibute `Groups` set to the static value `sac`.

Procedure

1. Go to *System* > *Administration* > *App Integration*.
2. In the *Trusted Identity Providers* section, click *Add a Trusted Identity Provider*.
3. In the dialog, enter the following properties:

Property	Description
Name	Enter a unique name, which will appear in the list of trusted identity providers.
Provider Name	Enter a unique name for the provide. This name can contain only alphabet characters (a-z & A-Z), numbers (0-9), underscore (_), dot (.), hyphen (-), and cannot exceed 36 characters.
Signing Certificate	Enter the signing certificate information for the third-party application server in X.509 Base64 encoded format.

4. Click *Add*.

   The identity provider is added to the list. Hover over it and select *Edit* to update it or *Delete* to delete it.

   You may need to use the *Authorization URL* and *Token URL* listed here to complete setup on your OAuth clients.

2.8 Delete Your Service Instance in SAP BTP

Delete your SAP Datasphere service instance in SAP BTP.

To do so, you must have SAP BTP administration authorization on the subaccount that is entitled to SAP Datasphere.

> **Note**
>
> When you delete your service instance, all your data will also be deleted. The tenant and all its data cannot be recovered.

1. In SAP BTP, select the subaccount and the space where the service instance was created.
2. Navigate to *Instances and Subscriptions*.
3. In the *Service Instances* page, find the SAP Datasphere service instance that you want to delete, click the button at the end of the row and select *Delete*, then click *Delete* in the confirmation dialog.
   You can view the progress of the deletion.

3 Managing Users and Roles

Create and manage users, manage secure access to SAP Datasphere using roles, and set up authentication for your users if you are using your own identity provider.

3.1 Configuring Identity Provider Settings

By default, SAP Cloud Identity Authentication is used by SAP Datasphere. We also support single sign-on (SSO), using your identity provider (IdP).

Related Information

Enable IdP-Initiated Single Sign On (SAP Data Center Only) [page 33]
Renewing the SAP Analytics Cloud SAML Signing Certificate [page 35]
Enabling a Custom SAML Identity Provider [page 36]

3.1.1 Enable IdP-Initiated Single Sign On (SAP Data Center Only)

By default, IdP-initiated SSO is not supported if SAP Datasphere is running on an SAP Data Center. To support IdP initiated SSO on an SAP Data Center, you must add a new assertion consumer service endpoint to your identity provider.

Prerequisites

SAP Datasphere can be hosted either on SAP data centers or on non-SAP data centers. Determine which environment SAP Datasphere is hosted in by inspecting your URL:

- A single-digit number, for example us1 or jp1, indicates an SAP data center.
- A two-digit number, for example eu10 or us30, indicates a non-SAP data center.

Procedure

1. Navigate to your IdP and find the page where you configure SAML 2.0 Single Sign On.
2. Find and copy your FQDN.

   For example, mysystem.wdf.sap-ag.de
3. Add a new assertion consumer service (ACS) endpoint that follows this pattern:

   ```
   https:// <FQDN>/
   ```

   For example, https://mysystem.wdf.sap-ag.de/
4. If you are using SAP Cloud Identity Authentication Service as your identity provider, the link to log onto SAP Datasphere through your identity provider will follow this pattern:

   ```
   https://<tenant_ID>.accounts.ondemand.com/saml2/idp/sso?
   sp=<sp_name>&index=<index_number>
   ```

   For example, https://testsystem.accounts999.ondemand.com/saml2/idp/sso?sp=mysystem.wdf.sap-ag.de.cloud&index=1

   > **Note**
   >
   > The pattern will vary depending on the identity provider you use.

   The following table lists the URL parameters you can use for IdP-initiated SSO.

Parameter	Mandatory	Description
sp	Yes	• This is the name of the SAML 2 service provider for which SSO is performed. • The sp_name value of the parameter equals the Entity ID of the service provider. • This parameter is needed for Identity Authentication to know which service provider to redirect the user to after successful authentication.
index	**Note** You can choose by the index the correct ACS endpoint for unsolicited SAML response processing. Provide the index parameter when the default ACS endpoint that has been configured via the administration console	• Enter the index number of the endpoint of the assertion consumer service of the service provider as the target of the SAML response. Otherwise, the identity provider uses the default endpoint configured for the trusted service provider. • If your IdP doesn't support indexing, you must choose

Parameter	Mandatory	Description
	cannot process unsolicited SAML responses.	between IdP-initiated SSO or SP-initiated SSO. You can either replace the default ACS endpoint to initiate an IdP SSO or continue using the default endpoint to initiate an SP SSO. • A non-digit value or a value for an index entry that is not configured returns an error message.

Results

Users will be able to use SAML SSO to log onto SAP Datasphere through their identity provider.

3.1.2 Renewing the SAP Analytics Cloud SAML Signing Certificate

To continue using SAML SSO, an administrator must renew the certificate before it expires.

Context

An email with details on how to renew the SAML X509 certificate is sent to administrators before the certificate expiry date. If the certificate expiry is less than 30 days away, a warning message appears when you log on to SAP Datasphere.

> **Note**
>
> If you click the *Renew* link on the warning message, you're taken to the *Security* tab on the (*Administration*) page.

Procedure

1. From the side navigation, go to (*System*) → (*Administration*)→ *Security*.
2. Select *Renew*.

   A confirmation dialog appears. When you confirm the renewal, a new metadata file is automatically downloaded.

> **Note**
>
> The renewal process takes around five minutes to complete.

3. If you use a custom identity provider, upload the SAP Datasphere metadata file to your SAML Identity Provider (IdP).

> **Note**
>
> This step is not required if you use SAP Cloud ID for authentication.

4. If you have live data connections to SAP HANA systems that use SAML SSO, you must also upload the new metadata file to your SAP HANA systems.
5. Log on to SAP Datasphere when five minutes has passed.

Results

If you are able to log on, the certificate renewal was successful. If you cannot logon, try one of the following troubleshooting tips.

If you use SAP Cloud ID for authentication:

1. Clear the browser cache.
2. Allow up to five minutes for the SAP Cloud ID service to switch to the new certificate.

If you use a custom identity provider for authentication:

1. Ensure the new metadata file has been uploaded to your IdP. For more information, see Enabling a Custom SAML Identity Provider [page 36].
2. Clear the browser cache.
3. Allow up to five minutes for your IdP to switch to the new certificate with the newly uploaded metadata.

3.1.3 Enabling a Custom SAML Identity Provider

By default, SAP Cloud Identity Authentication is used by SAP Datasphere. SAP Datasphere also supports single sign-on (SSO), using your identity provider (IdP).

Prerequisites

SAP Datasphere can be hosted on non-SAP data centers.

- You must have an IdP that supports SAML 2.0 protocol.
- You must be able to configure your IdP.
- You must be assigned to the *System Owner* role. For more information see Transfer the System Owner Role [page 97].

- If your users are connecting from Apple devices using the mobile app, the certificate used by your IdP must be compatible with Apple's App Transport Security (ATS) feature.

> **Note**
>
> A custom identity provider is a separate solution, like for example Azure AD, and is not part of SAP Analytics Cloud or SAP Datasphere. Therefore the change in configuration is to be applied directly in the solution, not within SAP Analytics Cloud or SAP Datasphere. Also no access to SAP Analytics Cloud or SAP Datasphere is required to make the change, only an access to the Identity Provider, eg Azure AD.

> **Note**
>
> Be aware that the SAML attributes for SAP Datasphere roles do not cover user assignment to spaces. A user who logs into a SAP Datasphere tenant through SSO must be assigned to the space in order to access the space. If you do not assign a user to a space, the user will not have access to any space.

Procedure

1. From the side navigation, go to (*System*) → (*Administration*) →*Security*.

   If you've provisioned SAP Datasphere prior to version 2021.03 you'll see a different UI and need go to (*Product Switch*) → (*Analytics*) → (*System*) → (*Administration*) → *Security*.
2. Select *(Edit)*.
3. In the *Authentication Method* area, select *SAML Single Sign-On (SSO)* if it is not already selected.

   > **Note**
   >
   > By default, SAP Cloud ID is used for authentication.
4. In *Step 1*, select *Download* and save the metadata file.
   A metadata file is saved.
5. Upload the metadata file to your SAML IdP.

   The file includes metadata for SAP Datasphere, and is used to create a trust relationship between your SAML Identity Provider and your SAP Datasphere system.
6. Optional: You can access the system from your SAML Identity Provider by adding a new assertion consumer service endpoint to your identity provider. For more information, see Enable IdP-Initiated Single Sign On (SAP Data Center Only) [page 33].
7. Map your SAML IdP user attributes and roles.

   If SAP Datasphere is running on an SAP data center, you must submit an SAP Product Support Incident using the component `LOD-ANA-ADM`. In the support ticket, indicate that you want to set up user profiles and role assignment based on custom SAML attributes, and include your SAP Datasphere tenant URL.

   > **Note**
   >
   > If SAP Datasphere is running on an SAP data center, and you want to continue using User Profiles and Role assignment using SAML attributes, you will need to open a support ticket each time you switch to a different custom IdP.

If SAP Datasphere is running on a non-SAP data center, you must configure your SAML IdP to map user attributes to the following case-sensitive allowlisted assertion attributes:

Attribute Name	Notes
email	Required if your `NameID` is "email".
Groups	Required. The value must be set to "sac", even in case of SAP Datasphere. The Groups attribute is a custom attribute and must be added if it does not exist yet. You need to contact your administrator to get the path where the mapping needs to be changed.
familyName	Optional. `familyName` is the user's last name (surname).
displayName	Optional.
functionalArea	Optional.
givenName	Optional. `givenName` is the user's first name.
preferredLanguage	Optional.
custom1	Optional. For SAML role assignment.
custom2	Optional. For SAML role assignment.
custom3	Optional. For SAML role assignment.
custom4	Optional. For SAML role assignment.
custom5	Optional. For SAML role assignment.

Example:

```
<AttributeStatement>
            <Attribute
                Name="email">
                <AttributeValue>abc.def@mycompany.com</AttributeValue>
            </Attribute>
            <Attribute
                Name="givenName">
                <AttributeValue>Abc</AttributeValue>
            </Attribute>
            <Attribute
                Name="familyName">
                <AttributeValue>Def</AttributeValue>
            </Attribute>
            <Attribute
                Name="displayName">
                <AttributeValue>Abc Def</AttributeValue>
            </Attribute>
            <Attribute
                Name="Groups">
                <AttributeValue>sac</AttributeValue>
            </Attribute>
            <Attribute
                Name="custom1">
                <AttributeValue>Domain Users</AttributeValue>
                <AttributeValue>Enterprise Admins</AttributeValue>
                <AttributeValue>Enterprise Key Admins</AttributeValue>
            </Attribute>
        </AttributeStatement>
```

> **Note**
>
> If you are using the SAP Cloud Identity Authentication service as your IdP, map the `Groups` attribute under *Default Attributes* for your SAP Datasphere application. The remaining attributes should be mapped under *Assertion Attributes* for your application.

8. Download metadata from your SAML IdP.
9. In *Step 2*, select *Upload*, and choose the metadata file you downloaded from your SAML IdP.
10. In *Step 3*, select a *User Attribute*.

    The attribute will be used to map users from your existing SAML user list to SAP Datasphere`NameID` used in your custom SAML assertion:

    ```
    <NameID Format="urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified"><Your Unique Identifier></NameID>
    ```

    Determine what your `NameID` maps to in your SAP Datasphere system. It should map to . The user attribute you select must match the *User ID*, *Email* or a custom attribute. You can view your SAP Datasphere user attributes in *Security* > *Users*.

    > **Note**
    >
    > `NameID` is case sensitive. The *User ID*, *Email*, or *Custom SAML User Mapping* must match the values in your SAML IdP exactly. For example, if the `NameId` returned by your SAML IdP is **`user@company.com`** and the email you used in SAP Datasphere is **`User@company.com`** the mapping will fail.

    Choose one of the following options:

    - *USER ID*: If `NameID` maps to the SAP Datasphere *User ID*.
    - *Email*: If `NameID` maps to SAP Datasphere *Email* address.

      > **Note**
      >
      > If your `NameID` email is not case-sensitive and contains mixed-case, for example `User@COMPANY.com`, consider choosing *Custom SAML User Mapping* instead.

    - *Custom SAML User Mapping*: If `NameID` maps to a custom value.

      > **Note**
      >
      > If you select this option, there will be a new column named *SAML User Mapping* in *Security* > *Users*. The. After switching to your SAML IdP, you must manually update this column for all existing users.

> **Note**
>
> If you are using a live connection to SAP S/4HANA Cloud Edition with *OAuth 2.0 SAML Bearer Assertion*, `NameId` must be identical to the user name of the business user on your SAP S/4HANA system.
>
> For example, if you want to map an SAP Datasphere user with the user ID **SACUSER** to your SAP S/4HANA Cloud user with the user name **S4HANAUSER**, you must select *Custom SAML User Mapping* and use **S4HANAUSER** as the *Login Credential* in Step 10.

> If you are using SAP Cloud Identity as your SAML IdP, you can choose *Login Name* as the *NameID* attribute for SAP Datasphere, then you can set the login name of your SAP Datasphere user as **S4HANAUSER**.

11. Optional: Enable *Dynamic User Creation*.

    When dynamic user creation is enabled, new users will be automatically created using the default role and will be able to use SAML SSO to log onto SAP Datasphere. After users are created, you can set roles using SAML attributes.

    > **Note**
    >
    > Automatic user deletion is not supported. If a user in SAP Datasphere is removed from your SAML IdP, you must go to *Security* > *Users* and manually delete users. For more information, see Deleting Users [page 51].
    >
    > If this option is enabled, dynamic user creation still occurs even when SAML user attributes have not been set for all IdP users. To prevent a user from being automatically created, your SAML IdP must deny the user access to SAP Datasphere.

12. In *Step 4*, enter `<Your Unique Identifier>`.

    This value must identify the system owner. The *Login Credential* provided here are automatically set for your user.

    > **Note**
    >
    > The *Login Credential* depends on the *User Attribute* you selected under *Step 3*.

13. Test the SAML IdP setup, by logging in with your IdP, and then clicking *Verify Account* to open a dialog for validation.

    In another browser, log on to the URL provided in the *Verify Your Account* dialog, using your SAML IdP credentials. You can copy the URL by selecting *(Copy)*.

    You must use a private session to log onto the URL; for example, guest mode in Chrome. This ensures that when you log on to the dialog and select SAP Datasphere, you are prompted to log in and do not reuse an existing browser session.

    > **Note**
    >
    > If SAP Datasphere is running on a non-SAP data center, upon starting the verification step, you will see a new screen when logging into SAP Datasphere. Two links will be displayed on this page. One will link to your current IdP and the other will link to the new IdP you will switch to. To perform the *Verify Account* step, use the link for the new IdP. Other SAP Datasphere users can continue logging on with the current IdP. Once you have completed Step 16 and the IdP switch has completed, this screen will no longer appear.

    If you can log on successfully, the SAML IdP setup is correct.

14. In the *Verify Your Account* dialog, select *Check Verification*.
    If the verification was successful, a green border should appear around the *Login Credential* box.

15. Select *(Save)*.
    The *Convert to SAML Single Sign-On* confirmation dialog will appear.

16. Select *Convert*.

When conversion is complete, you will be logged out and directed to the logon page of your SAML IdP.

17. Log on to SAP Datasphere with the credentials you used for the verification step.
18. From the side navigation, go to (*Security*) → and (*Users*)look for the column of the *User Attribute* you selected in step 8.

    The values in this column should be a case sensitive match with the `NameId` sent by your IdP's SAML assertion.

    > **Note**
    >
    > If you selected *Custom SAML User Mapping* as *User Attribute*, you must manually update all fields in the *SAML User Mapping* column.

Results

Users will be able to use SAML SSO to log onto SAP Datasphere.

> **Note**
>
> You can also set up your IdP with your Public Key Infrastructure (PKI) so that you can automatically log in your users with a client side X.509 certificate.

Next Steps

Switch to a Different Custom IdP

If SAML SSO is enabled and you would like to switch to a different SAML IdP, you can repeat the above steps using the new SAML IdP metadata.

3.1.3.1 Disabling SAML SSO

You can revert your system to the default identity provider (SAP Cloud Identity) and disable your custom SAML IdP.

Procedure

1. From the side navigation, go to (*System*)→ (*Administration*) → *Security*.

   If you've provisioned SAP Datasphere prior to version 2021.03 you'll see a different UI and need go to (*My Products*) → (*Analytics*) → (*System*) → (*Administration*) → *Security*.

2. Select *(Edit)* .
3. In the *Authentication Method* area, select *SAP Cloud Identity (default)*.
4. Select *(Save)* .

Results

When conversion is complete, you will be logged out and directed to the SAP Cloud Identity logon page.

3.1.3.2 Updating the SAML IdP Signing Certificate

You can update the SAML identity provider (IdP) signing certificate.

Prerequisites

- You must have the metadata file that contains the new certificate from your custom IdP, and you must be logged into SAP Datasphere before your IdP switches over to using the new certificate.
- You must be the System Owner in SAP Datasphere.

Context

To upload the new metadata file, do the following:

Procedure

1. From the side navigation, go to (*System*) → (*Administration*) →*Security* .

   If you've provisioned SAP Datasphere prior to version 2020.03 you'll see a different UI and need go to (*My Products*) → (*Analytics*) → (*System*) → (*Administration*) → *Security*.
2. Select *(Edit)*
3. Under *Step 2*, select *Update* and provide the new metadata file.
4. Select *(Save)* and confirm the change to complete the update.

   The update will take effect within two minutes.

Results

> **Note**
>
> You do not have to redo *Step 3* or *Step 4* on the *Security* tab.

3.1.3.3 Identity Provider Administration

The Identity Provider Administration tool allows system owners to manage the custom identity provider configured with SAP Datasphere. Through the tool, the system owner can choose to upload new metadata for the current custom identity provider, or revert to using the default identity provider.

Prerequisites

- SAP Datasphere must already be configured to use a custom identity provider.
- You must be the system owner.

Procedure

1. Access the *Identity Provider Administration* tool using the following URL pattern:
   `https://console.<data center>.sapanalytics.cloud/idp-admin/`

   For example, if your SAP Datasphere system is on eu10, then the URL is:

   `https://console.eu10.sapanalytics.cloud/idp-admin/`

   If your SAP Datasphere system is on cn1, then the URL is:

   `https://console.cn1.sapanalyticscloud.cn/idp-admin/`

   If your tenant is on EUDP:

https://console-eudp.eu1.sapanalytics.cloud/idp-admin/
https://console-eudp.eu2.sapanalytics.cloud/idp-admin/

2. Log in with an S-user that has the same email address as the system owner of your system. If you don't yet have such an S-user, you can click the “Register” button and create a P-user.

   If you create a new P-user, you'll receive an email with an activation link that will let you set your password.
3. Once you're logged in, you'll see a list of SAP Datasphere systems for which you are the system owner.

   Select the system you want to work on by clicking on its row.

Once you're in the settings page for your system, you can see information about your current custom identity provider. If you need to reacquire your system's metadata, you can click the “Service Provider Metadata Download” link.

If you don't want to manage your custom identity provider through Identity Provider Administration, you can disconnect your system by clicking “Disconnect IdP Admin from your system”.

4. To proceed with either reverting to the default identity provider or updating the current custom identity provider, select the corresponding radio button and then click “Step 2”.

> ### Note
>
> Your SAP Datasphere system is connected to the Identity Provider Administration tool by default. The connection status for your system is displayed under the “Status” column of the systems list page. If you'd like to disconnect your system from the console, you can do so in either of two places:
>
> - In SAP Datasphere, navigate to *System* > *Administration* > *Security* > *Optional: Configure Identity Provider Administration Tool*, click the *Connected* switch, and then save the changes.
> - Click “Disconnect IdP Admin from your system” after selecting your system in Identity Provider Administration.

5. (Optional) Revert to the default identity provider.

   Choose this option if you're having problems logging in with your custom identity provider and would like to revert to the default identity provider. Once the reversion has finished, you can exit the Identity Provider Administration tool and log in to your SAP Datasphere system to reconfigure your custom identity provider.

   a. Select the “Yes” radio button to revert to the default IdP.
   b. Select “Yes” in the confirmation dialog to revert your authentication method back to the default IdP.
   c. Click “Step 3” to proceed to the validation step.
   d. Click “Log into SAP Datasphere” to open a new tab and navigate to your system. Log in with your default identity provider credentials. If you get an error saying “Your profile is not configured”, please create a support ticket under the component LOD-ANA-BI.

6. (Optional) Upload new metadata for the current custom identity provider.

   Choose this option if you need to reconfigure trust between your custom identity provider and your SAP Datasphere system. A common use case is to upload new metadata from your identity provider when a new signing certificate has been generated.

   a. Click “Browse” to select the new metadata file for your current custom identity provider.
   b. Click “Upload File” to upload the provided metadata file. After the upload is successful, it can take up to five minutes for the new metadata file to be applied.
   c. Click “Step 3” to proceed to the validation step.
   d. Click “Log into SAP Datasphere” to open a new tab and navigate to your system. If you have any login problems related to the identity provider configuration, as opposed to a user-specific problem, you can return to the Identity Provider Administration tool and either re-upload the metadata file or revert to the default identity provider.

3.2 Managing SAP Datasphere Users

You can create and modify users in SAP Datasphere in several different ways.

Creating Users

You can create users in the following ways:

Method	More Information
Create individual users in the *Users* list	Creating a New User [page 45]
Import multiple users from a CSV file	Importing or Modifying Users from a File [page 47]

Modifying Users

You can modify existing users in the following ways:

Modification	More Information
Export user data to a CSV file, to synchronize with other systems	Exporting Users [page 49]
Update the email address a user logs on with	Updating User Email Addresses [page 50]
Delete users	Deleting Users [page 51]

3.2.1 Creating a New User

You can create individual users in SAP Datasphere.

Prerequisites

You can select one or more roles while you're creating the user. Before getting started creating users, you might want to become familiar with the global roles and scoped roles. You can still assign roles after you've created the users.

Type of Role	Description	More Information
Global Roles	A role that enables users assigned to it to perform actions that are not space-related, typically a role that enables to administrate the tenant. A standard or custom role is considered as global when it includes global privileges.	Managing Roles and Privileges [page 52]
Scoped Roles	A role that inherits a set of scoped privileges from a standard or custom role and grants these privileges to users for use in the assigned spaces.	Create a Scoped Role to Assign Privileges to Users in Spaces [page 57]

Context

The method described here assumes that SAP Datasphere is using its default authentication provider. If you are using a custom SAML Identity Provider, you must provide slightly different information, depending upon how your SAML authentication is configured.

Procedure

1. Go to ☰ (*Expand*) > 🔒 (*Security*) > 👤 (*Users*).
2. Select + *(New)* to add a new user to the user management table.
3. Enter a *User ID*.

   Each user needs a unique ID. Only alphanumeric and underscore characters are allowed. The maximum length is 127 characters.
4. Enter the user name details.

   Only *Last Name* is mandatory, but it is recommended that you provide a *First Name*, *Last Name*, and *Display Name*. *Display Name* will appear in the screens.
5. Enter an *Email* address.

   A welcome email with logon information will be sent to this address.

   > ### Note
   >
   > The *Manager* column is not relevant for SAP Datasphere users.
6. Select the icon and choose one or more roles from the list.
7. Select 💾 *(Save)*.

Results

A welcome email including an account activation URL will be sent to the user, so that the user can set an initial password and access the system.

> **Note**
>
> In addition to the standard workflows, you can also create users via the command line (see Manage Users via the Command Line).

3.2.2 Importing or Modifying Users from a File

You can create new users or batch-update existing users by importing user data that you have saved in a CSV file.

Prerequisites

The user data you want to import must be stored in a CSV file. At minimum, your CSV file needs columns for `UserID`, `LastName`, and `Email`, but it is recommended that you also include `FirstName` and `DisplayName`.

If you want to assign new users different roles, include a `Roles` column in the CSV file. The role IDs used for role assignment are outlined in Standard Roles Delivered with SAP Datasphere [page 54].

For existing users that you want to modify, you can create the CSV file by first exporting a CSV file from SAP Datasphere. For more information, see Exporting Users [page 49].

> **Note**
>
> The first name, last name, and display name are linked to the identity provider, and can't be changed in the User list page, or when importing a CSV file. (In the User list page, those columns are grayed out.)
>
> To edit those values, you'll need to use the user login, and edit that user's profile.

Edit the downloaded CSV file to remove columns whose values you don't want to modify, and to remove rows for users whose values you don't want to modify. Do not modify the `USERID` column. This ensures that entries can be matched to existing users when you re-import the CSV.

These are the available mapping parameters when importing CSV user data:

Parameter	Description
User ID	
First Name	
Last Name	
Display Name	

Parameter	Description
Email	
Manager	
Roles	
Mobile	
Phone	
Office Location	
Function Area	Can be used to refer to a user's team or area within their organization.
Job Title	
Clean up notifications older than	Set in user settings: when to automatically delete notifications.
Email Notification	Set in user settings.
Welcome message	Message that is shown to the user on the home screen.
Page tips	Enabled/disabled via the help center (deprecated).
Closed Page tips	Closed page tips are tracked so that they are not shown again.
Closed Item Picker Tips	Closed tooltips are tracked so that they won't be reopened again (for first time users).
Current Banner	Saves which banner is currently showing.
Last Banner	The UUID of the last closed banner.
Last Maintenance Banner Version	The version when the last maintenance banner was shown.
Marketing email opt in	Set in user settings.
Homescreen content is initialized	If default tiles have been set for the home screen.
Expand Story Toolbar	Set in user settings.
Is user concurrent	If the user has a concurrent license.
On the *Edit Home Screen* dialog, a user can override all the default preferences that have been set by the administrator for the system (*System* > *Administration* > *Default Appearance*). These are the preferences:	
Override Background Option	
Override Logo Option	
Override Welcome Message	
Override Home Search To Insight	
Override Get Started	
Override Recent Stories	
Override Recent Presentations	
Override Calendar Highlights	

Procedure

1. Go to ☰ (*Expand*) ▸ 🔒 (*Security*) ▸ 👤 (*Users*).
2. Select *(Import Users)* ▸ *Import Users from File*.
3. In the *Import Users* dialog, choose *Select Source File* to upload your CSV file.
4. Choose *Create Mapping* to assign the fields of your user data from the CSV file to the fields in user management.
5. Select the appropriate entries for the *Header*, *Line Separator*, *Delimiter*, and *Text Qualifier*.
6. Select *OK* when you've finished mapping.
7. In the *Import Users* dialog, choose *Import* to upload your CSV file according to the defined mapping.

3.2.3 Exporting Users

If you want to synchronize SAP Datasphere user data with other systems, you can export the data to a CSV file.

Procedure

On the *Users* page of the *Security* area, choose *(Export)*.

Results

The system exports all user data into a CSV file that is automatically downloaded to your browser's default download folder.

The CSV file contains these columns:

Column	Description
USER_NAME	
FIRST_NAME	
LAST_NAME	
DISPLAY_NAME	
EMAIL	
MANAGER	
ROLES	Roles assigned to the user.
SAML_USER_MAPPING	SAML property for the user (if SAML enabled).
MOBILE	Set in user preferences.

Column	Description
OFFICE_PHONE	Set in user preferences.
OFFICE_ADDRESS	Set in user preferences.
AGILE_BI_ENABLED_BY_DEFAULT	Opt in for the agile data preparation feature.
JOB_TITLE	Set in user preferences.
MARKETING_EMAIL_OPT_IN	Set in user preferences.
IS_CONCURRENT	Licensing attribute to indicate whether the user is consuming a named licensed user account (0) or a concurrent licensed user account (1).
DEFAULT_APP	The application that will launch when you access your SAP Datasphere URL. The default application can be set in *System* > *Administration* > *System Configuration* or in the user settings.
On the *Edit Home Screen* dialog, a user can override all the default preferences that have been set by the administrator for the system (*System* > *Administration* > *Default Appearance*). These are the preferences:	
OVERRIDE_BACKGROUND_OPTION	
OVERRIDE_LOGO_OPTION	
OVERRIDE_WELCOME_MESSAGE_FLAG	
OVERRIDE_HOME_SEARCH_TO_INSIGHT_FLAG	
OVERRIDE_GET_STARTED_FLAG	
OVERRIDE_RECENT_FILES_FLAG	
OVERRIDE_RECENT_STORIES_FLAGOVERRIDE_RECENT_STORIES_FLAG	
OVERRIDE_RECENT_PRESENTATIONS_FLAG	
OVERRIDE_RECENT_APPLICATIONS_FLAG	
OVERRIDE_CALENDAR_FLAG	
OVERRIDE_FEATURED_FILES_FLAG	

3.2.4 Updating User Email Addresses

You can update the user email addresses used for logon.

When you create a user, you must add an email address. The email address is used to send logon information.

To edit a user's email address, go to the *Users* page of the *Security* area, and select the email address you want to modify. Add a new email address and press *Enter*, or select another cell to set the new address.

If the email address is already assigned to another user, a warning will appear and you must enter a new address. Every user must be assigned a unique email address.

A new logon email will be sent to the updated address.

As long as a user has not logged on to the system with the new email address, the email address will appear in a pending state on the *Users* list.

Related Information

Creating a New User [page 45]
Importing or Modifying Users from a File [page 47]

3.2.5 Deleting Users

You can delete users.

Procedure

1. In the *Users* management table, select the user ID you want to delete by clicking the user number in the leftmost column of the table.
   The whole row is selected.
2. Choose *(Delete)* from the toolbar.
3. Select *OK* to continue and remove the user from the system.

Related Information

Creating a New User [page 45]
Importing or Modifying Users from a File [page 47]
Updating User Email Addresses [page 50]

3.2.6 Set a Password Policy for Database Users

Users with the *DW Administrator* role (administrators) can set a password policy to cause database user passwords to expire after a specified number of days.

Context

Users with the *DW Space Administrator* role (space administrators) can create database users in their spaces to allow the connection of ETL tools to write to and read from Open SQL schemas attached to the space schema (see Integrating Data via Database Users/Open SQL Schemas).

Procedure

1. In the side navigation area, click (*System*) (*Configuration*) *Security*.
2. In the *Password Policy Configuration* section, enter the number of days after which a database user's password will expire.

   After this period, the user will be prompted to set a new password.

   > **Note**
   >
   > The password policy applies only to database users where the *Enable Password Policy* property is selected, for both existing and new users. If a user does not log on with their initial password during this period, they will be deactivated until their password is reset.

3.3 Managing Roles and Privileges

Assigning roles to your users maintains access rights and secures your information in SAP Datasphere.

A role is a set of privileges and permissions.

SAP Datasphere delivers a set of standard roles and you can create your own custom roles:

- Standard role - A role delivered with SAP Datasphere that includes a set of privileges. As a best practice, a tenant administrator can use these roles as templates for creating custom roles for different business needs. See Standard Roles Delivered with SAP Datasphere [page 54].
- Custom role - A role that a tenant administrator creates to choose specific privileges as needed. See Create a Custom Role [page 56].

Each standard or custom role is either a global role or a template for scoped roles:

- Global role - A role that enables users assigned to it to perform actions that are not space-related, typically a role that enables to administrate the tenant. A standard or custom role is considered as global when it

includes global privileges. A tenant administrator can assign a global role to the relevant users. See Assign Users to a Role [page 63].

- Scoped role - A role that inherits a set of privileges from a standard or custom role and assigns them to one or more users for one or more spaces. Users assigned to a scoped role can perform actions in the assigned spaces. A tenant administrator can create a scoped role. See Create a Scoped Role to Assign Privileges to Users in Spaces [page 57].

For more information on global and scoped privileges, see Privileges and Permissions [page 65].

> ⚠ Caution
>
> Scoped roles and all related features will be rolled out to all tenants over the course of a number of versions. For more details, see SAP Note 3380409.

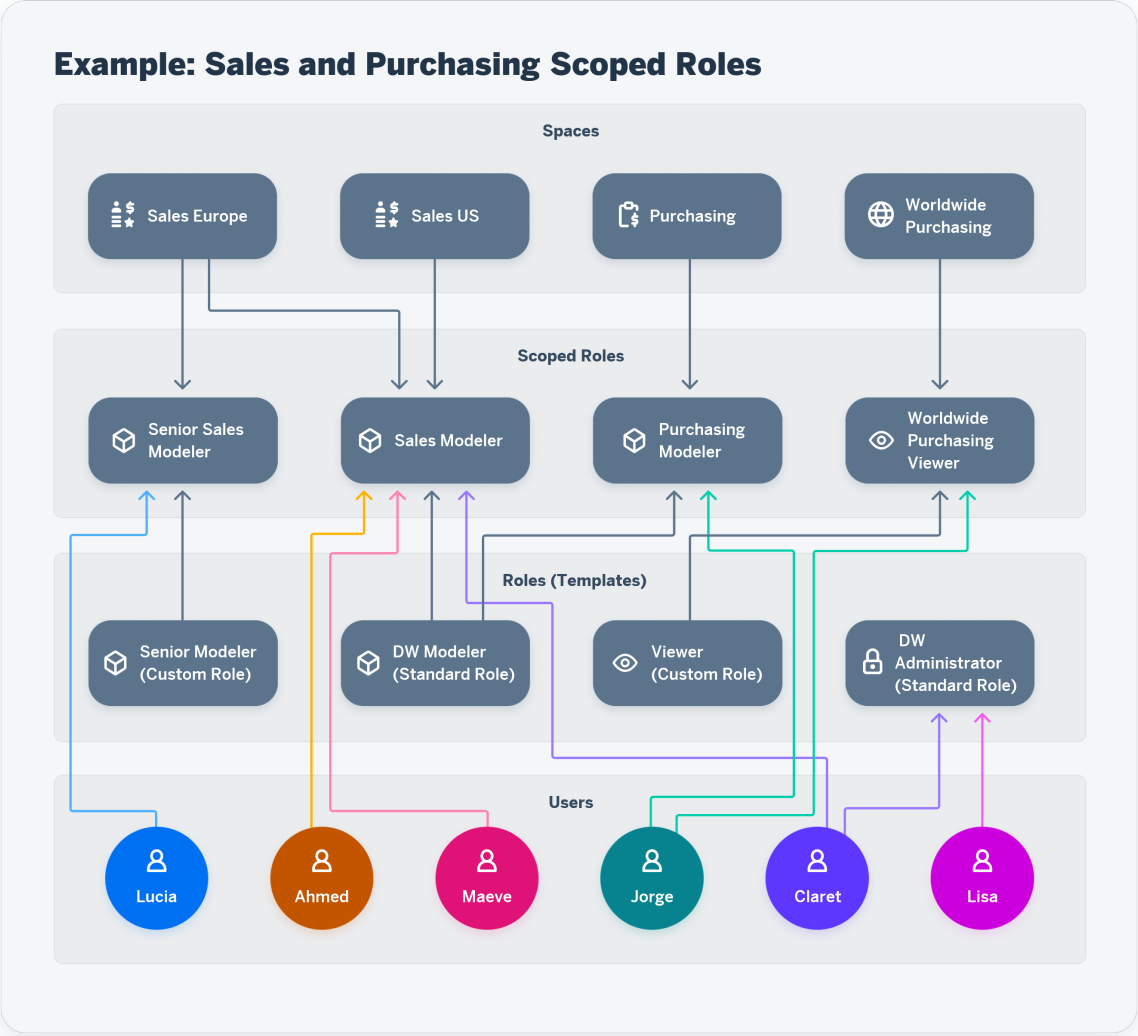

Users have relevant privileges depending on which actions they can do in the spaces.

- Lisa administers the SAP Datasphere tenant.
- Claret administers the SAP Datasphere tenant and also has modeler privileges in the two spaces Sales Europe and Sales US.

- Jorge has purchasing modeler privileges in the Purchasing space and has viewer privileges in the Worldwide Purchasing space.
- Maeve and Ahmed have modeler privileges in the two spaces Sales Europe and Sales US.
- Lucia has modeler privileges in the Sales Europe space.

3.3.1 Standard Roles Delivered with SAP Datasphere

SAP Datasphere is delivered with several standard roles. A standard role includes a predefined set of privileges and permissions.

A DW Administrator can use standard roles as templates for creating custom roles with a different set of privileges (see Create a Custom Role [page 56]). You can also use the standard roles that include scoped privileges as templates for creating scoped roles (see Create a Scoped Role to Assign Privileges to Users in Spaces [page 57]). You can assign the standard roles that contain global privileges (DW Administrator, Catalog Administrator and Catalog User) directly to users.

> **Note**
>
> You cannot delete nor edit standard roles.

> **Caution**
>
> Scoped roles and all related features will be rolled out to all tenants over the course of a number of versions. For more details, see SAP Note 3380409.

In the side navigation area, click (*Security*) (*Roles*). The following standard roles are available:

- Roles providing privileges to administer the SAP Datasphere tenant:
  - **System Owner** - Includes all user privileges to allow unrestricted access to all areas of the application. Exactly one user must be assigned to this role.
  - **DW Administrator** - Can create users, roles and spaces and has other administration privileges across the SAP Datasphere tenant. Cannot access any of the apps (such as the *Data Builder*).
- Roles providing privileges to work in SAP Datasphere spaces:
  - **DW Space Administrator** (template) - Can manage all aspects of the spaces users are assigned to (except the *Space Storage* and *Workload Management* properties) and can create data access controls.
    - *DW Scoped Space Administrator* - This predefined scoped role is based on the DW Space Administrator role and inherits its privileges and permissions.

      > **Note**
      >
      > Users who are space administrators primarily need scoped permissions to work with spaces, but they also need some global permissions (such as Lifecycle when transporting content packages). To provide such users with the full set of permissions they need, they must be assigned to a scoped role (such as the *DW Scoped Space Administrator*) to receive the necessary scoped privileges, but they also need to be assigned directly to the *DW Space Administrator* role (or a custom role that is based on the *DW Space Administrator* role) in order to receive the additional global privileges.

    - *DW Integrator* (template) - Can integrate data via connections and can manage and monitor data integration in a space.
        - *DW Scoped Integrator* - This predefined scoped role is based on the DW Integrator role and inherits its privileges and permissions.
    - **DW Modeler** (template) - Can create and edit objects in the *Data Builder* and *Business Builder* and view data in objects.
        - *DW Scoped Modeler* - This predefined scoped role is based on the DW Modeler role and inherits its privileges and permissions.
    - **DW Viewer** (template) - Can view objects and view data output by views that are exposed for consumption in spaces.
        - *DW Scoped Viewer* - This predefined scoped role is based on the DW Viewer role and inherits its privileges and permissions.
- Roles providing privileges to consume the data exposed by SAP Datasphere spaces:
    - **DW Consumer** (template) - Can consume data exposed by SAP Datasphere spaces, using SAP Analytics Cloud, and other clients, tools, and apps. Users with this role cannot log into SAP Datasphere. It is intended for business analysts and other users who use SAP Datasphere data to drive their visualizations, but who have no need to access the modeling environment.
        - *DW Scoped Consumer* - This predefined scoped role is based on the DW Consumer role and inherits its privileges and permissions.
- Roles providing privileges to work in the SAP Datasphere catalog:
    - **Catalog Administrator** - Can set up and implement data governance using the catalog. This includes connecting the catalog to source systems for extracting metadata, building business glossaries, creating tags for classification, and publishing enriched catalog assets so all catalog users can find and use them. Must be used in combination with another role such as *DW Viewer* or *DW Modeler* for the user to have access to SAP Datasphere.
    - **Catalog User** - Can search and discover data and analytics content in the catalog for consumption. These users may be modelers who want to build additional content based on official, governed assets in the catalog, or viewers who just want to view these assets. Must be used in combination with another role such as *DW Viewer* or *DW Modeler* for the user to have access to SAP Datasphere.

> **Note**
>
> Users are assigned roles in particular spaces via scoped roles. One user may have different roles in different spaces depending on the scoped role they're assigned to. See Create a Scoped Role to Assign Privileges to Users in Spaces [page 57].

> **Note**
>
> Please note for SAP Datasphere tenants that were initially provisioned prior to version 2021.03, you need the following additional roles to work with stories:
>
> - BI Content Creator - Creates and changes stories.
> - BI Content Viewer - Views stories.

Roles and Licenses

The standard roles are grouped by the license type they consume and each user's license consumption is determined solely by the roles that they've been assigned. For example, a user who has been assigned only the *DW Administrator* standard role consumes only a *SAP Datasphere* license.

Planning Professional, Planning Standard as well as Analytics Hub are SAP Analytics Cloud specific license types. For more information, see Understand Licenses, Roles, and Permissions in the *SAP Analytics Cloud* documentation.

3.3.2 Create a Custom Role

You can create a custom role using either a blank template or a standard role template and choosing privileges and permissions as needed.

Prerequisites

To create a custom role, you need the DW Administrator role.

Context

You can create a custom role to enable users to do either global actions on the tenant or actions that are specific to spaces.

- If you create a custom role for global purposes, you should include only global privileges and permissions. You can then assign the role to the relevant users.
- If you create a custom role for space-related purposes, you should include only scoped privileges and permissions. As a second step, you need to create a scoped role based on this custom role to assign users and spaces to the set of privileges included. See Create a Scoped Role to Assign Privileges to Users in Spaces [page 57].

You should not mix global and scoped privileges in a custom role.

- If you include a scoped privilege in a custom role that you create for global purposes, the privilege is ignored.
- If you include a global privilege in a custom role that you want to use as a template for a scoped role, the privilege is ignored .

> **Note**
>
> Some users, such as space administrators, primarily need scoped permissions to work with spaces, but they also need some global permissions (such as Lifecycle when transporting content packages). To provide such users with the full set of permissions they need, you can include both the relevant global privileges and scoped privileges in the custom role you will use as a template for the scoped role. Each

space administrator is then assigned to the scoped role to receive the necessary scoped privileges, but they are also assigned directly to the custom role in order to receive the additional global privileges.

For more details about global and scoped privileges, see Privileges and Permissions [page 65].

> ⚠ Caution
>
> Scoped roles and all related features are rolled out to all tenants over the course of a number of versions. For more details, see SAP Note 3380409.

Procedure

1. Go to ☰ (*Expand*) > 🔒 (*Security*) > 🔧 (*Roles*).
2. To create a custom role, click ＋ *(Add Role)* and select *Create a Custom Role*.
3. Enter a unique name for the role and select the license type SAP Datasphere.
4. Select *Create*.
5. Select a role template.

   The role templates are the predefined standard roles associated with the SAP Datasphere license type. If you wish to create a role without extending a predefined standard role, choose the blank template. After you select a template, a page opens showing you the individual permissions assigned to the privileges that have been defined for the role template you chose.
6. Select the permissions for your new role for every privilege type. The permission privileges represent an area, app, or tool in SAP Datasphere while the permissions (create, read, update, delete, execute, maintain, share, and manage) represent the actions a user can perform. For more details about global and scoped privileges, see Privileges and Permissions [page 65].
7. If you want to change the role template that your new custom role will be based on, select (*Select Template*), and choose a role.
8. Save your new custom role.

   > ⓘ Note
   >
   > You can assign the role to a user from the *Users* page or - only if you've created a custom role for global purposes (and not for space-related purposes) - from the *Roles* page. Whether you create users first or roles first does not matter. See Assign Users to a Role [page 63].

3.3.3 Create a Scoped Role to Assign Privileges to Users in Spaces

A scoped role inherits a set of scoped privileges from a standard or custom role and grants these privileges to users for use in the assigned spaces.

This topic contains the following sections:

- Introduction to Scoped Roles [page 58]
- Create a Scoped Role [page 60]
- Add Spaces to a Scoped Role [page 61]
- Remove Spaces from a Scoped Role [page 61]
- Add Users to a Scoped Role [page 62]
- Remove Users from a Scoped Role [page 62]

> ⚠ Caution
>
> Scoped roles and all related features will be rolled out to all tenants over the course of a number of versions. For more details, see SAP Note 3380409.

Introduction to Scoped Roles

A user with the DW Administrator role can create scoped roles.

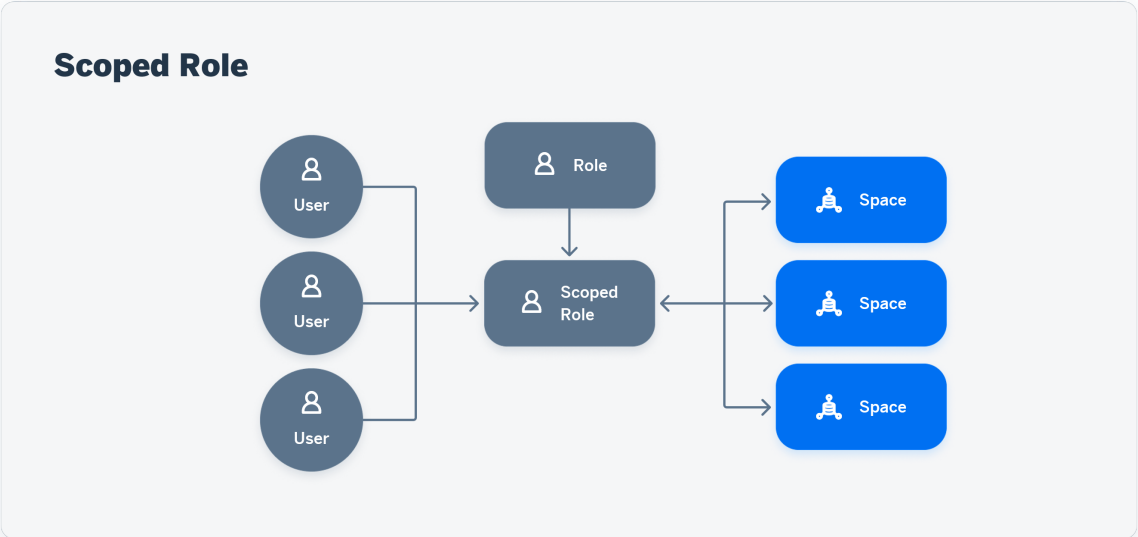

A DW Administrator can assign a role to multiple users in multiple spaces, in a single scoped role. As a consequence, a user can have different roles in different spaces: be a modeler in space Sales Germany and Sales France and a viewer in space Europe Sales.

You can create a scoped role based on a standard role or on a custom role. In both cases, the scoped role inherits the privileges from the standard or custom role. You cannot edit the privileges of a scoped role or of a standard role. You can edit the privileges of a custom role. To create a scoped role with a different set of privileges, create a custom role with the set of privileges wanted and then create the scoped role from the custom role. You can then change the privileges of the custom role as needed, which will also change the privileges of all the scoped roles that are based on the custom role.

Users who are granted the DW Space Administrator role via a scoped role can add or remove users to or from their spaces and the changes are reflected in the scoped roles. See Control User Access to Your Space.

We recommend that you create scoped roles by logical groups of spaces.

In the following example, the DW administrator begins assigning users to the three Sales spaces by creating the appropriate scoped roles:

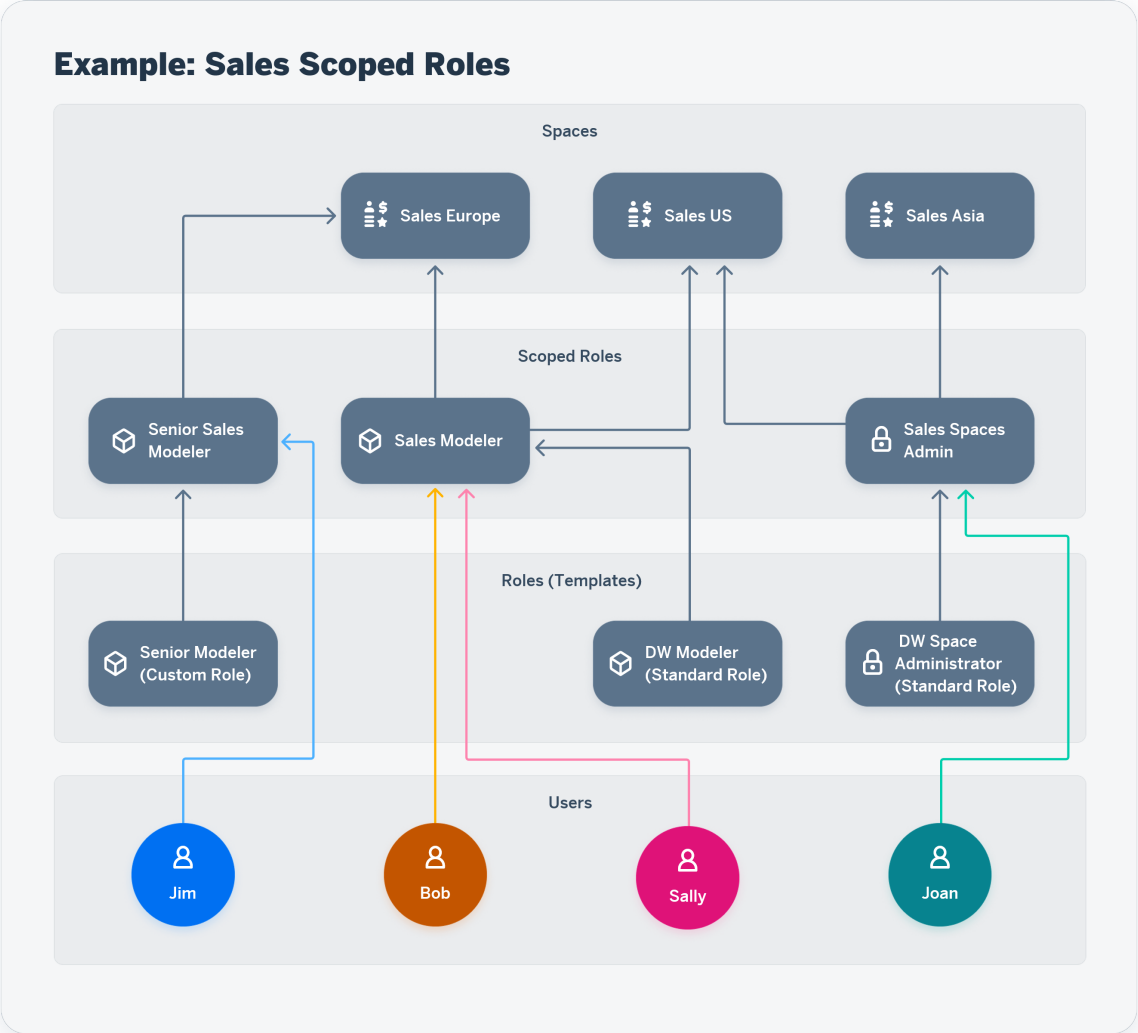

She creates three scoped roles based on standard and custom roles and assigns the users to the spaces as follows:

Scoped Roles	Roles (Templates)	Users	Spaces
Sales Modeler	DW Modeler standard role	Sally	Sales Europe
		Bob	Sales US

Scoped Roles	Roles (Templates)	Users	Spaces
Senior Sales Modeler	Custom role "Senior Modeler" based on the DW Modeler standard role + these privileges (and permissions): • Data Warehouse Data Integration (Execute) • Data Warehouse Remote Connection (Create, Read, Update and Delete)	Jim	Sales Europe
Sales Spaces Admin	DW Space Administrator standard role + this privilege (permission): • Scoped Role User Assignment (Manage)	Joan	Sales US Sales Asia

If Bob no longer needs to work in the space Sales US, the DW administrator can unassign Bob from Sales US in the scoped role Sales Modeler.

As Joan has the role of space administrator for the space Sales US, she can also unassign Bob from Sales US directly in the space page (in the *Space Management*). The user assignment change is automatically reflected in the Sales Modeler scoped role.

Later on, Bob needs the space administration privileges for the space Sales Asia. From the page of the space Sales Asia, Joan assigns Bob to the space with the Sales Space Admin scoped role.

For more information on scoped roles, see the blog Preliminary Information SAP Datasphere– Scoped Roles (published in September 2023).

Create a Scoped Role

> **Note**
>
> In addition to the standard workflows, you can also create scoped roles and assign scopes and users to them via the command line (see Manage Scoped Roles via the Command Line).

1. In the side navigation area, click (*Security*) > (*Roles*) and click your scoped role to open it.
2. Click + *(Add Role)* and select *Create a Scoped Role*.

   > **Note**
   >
   > As an alternative to creating a scoped role, you can use one of the predefined scoped roles that are delivered with SAP Datasphere in the *Roles* page and directly assign spaces and users to them.

3. Enter a unique name for the role and select the license type SAP Datasphere.

4. Click *Create*.
5. Select the role template, which can either be a standard role template or a custom role and click *Save*.
6. As your scoped role inherits privileges from the template you've chosen, you cannot edit the privileges, except for the one privilege *Scoped Role User Assignment* (Manage). If you're creating a scoped role for space administration purposes, you should select this privilege that allows to manage user assignment in a space.

You can then assign spaces and users to the new scoped role. The spaces and users must be created beforehand and you must assign spaces before assigning users to them.

> **Note**
>
> If you're creating a scoped role to assign space administration privileges to certain users in certain spaces, you can either do as follows:
>
> - Create a scoped role based on the standard role template DW Space Administrator and, to allow user assignment, select the privilege (permission) *Scoped Role User Assignment* privilege (Manage), which is the only privilege you can select, as the rest of the privileges are inherited from the template. Then, assign one or more spaces and one or more users to the spaces.
> - Open the predefined scoped role DW Scoped Space Administrator and assign one or more spaces and one or more users to the spaces. *Scoped Role User Assignment* (Manage) is selected by default.
>
> The users can manage the spaces they're assigned to.

Add Spaces to a Scoped Role

To add spaces to a scoped role, the spaces must be created beforehand.

1. In the side navigation area, click (*Security*) (*Roles*) and click your scoped role to open it.
2. Click *Add Scopes*, select one or more spaces in the dialog *Add Scopes* and click *Save*.

> **Note**
>
> By default, all users of the scoped role are automatically assigned to the spaces you've just added. You can change this and assign only certain members to certain spaces in the *Users* page of the scoped role.

Remove Spaces from a Scoped Role

1. In the side navigation area, click (*Security*) (*Roles*) and click your scoped role to open it.
2. Click *[number] Scopes*.
3. In the *Selected Scopes* area of the dialog *Add Scopes*, click the cross icon for each space that you want to remove from the role, then click *Save*.
   All users that were assigned to the spaces you've just removed are automatically removed from the scoped role.

Add Users to a Scoped Role

To add users to a scoped role, the users must be created beforehand.

1. In the side navigation area, click (*Security*) (*Roles*) and click your scoped role to open it.
2. Click *Users*. All user assignements are displayed in the *Users* page.
   - To individually select users and assign them to spaces, click + *(Add Users to Scopes)*, then *Add New Users to Scopes*. Select one or more users in the wizard *Add Users to Scopes* and click *Next Step*.

     > **Note**
     >
     > By default, the added users are automatically assigned to all the spaces included in the scoped role. If you want to modify this, select the one or more spaces to which you want to assign the users.

     Click *Next Step* and *Save*.

     > **Note**
     >
     > You can also add a user to a scoped role from the (*Users*) area. In such a case, the user is automatically assigned to all the spaces included in the scoped role. See Assign Users to a Role [page 63].

   - To assign all users included in the scoped role to one or more spaces. To do so, click + *(Add Users to Scopes)*, then *Add All Current Users to Scopes*. Select one or more spaces in the wizard *Add Users to Scopes* and click *Next Step* and *Save*.
   - To assign all users of the tenant to one or more spaces, click + *(Add Users to Scopes)*, then *Add All Users to Scopes*. Select one or more spaces in the wizard *Add Users to Scopes* and click *Next Step* and *Save*.

> **Restriction**
>
> A user can be assigned to a maximum of 100 spaces across all scoped roles.

> **Note**
>
> In the *Users* page, you can filter users and spaces to see for example to which spaces and roles a user is assigned to.

Once you've assigned a user to a space with the DW Space Administrator role via a scoped role, this user can manage the users for its space directly in the page of its space (in the *Space Management*). See Control User Access to Your Space.

Remove Users from a Scoped Role

1. In the side navigation area, click (*Security*) (*Roles*) and click your scoped role to open it.
2. Click *Users*. All user assignements are displayed in the *Users* page.
3. Check the relevant rows (a row corresponding to a combination of one user and one space) and click the garbage icon. The users cannot access the spaces they were previously assigned to in the scoped role.

3.3.4 Assign Users to a Role

There are multiple ways to assign users to a role.

Prerequisites

To assign roles, you need the DW Administrator role.

Assign or Update an Individual User's Role

1. In the side navigation area, click (*Security*) > (*Users*).
2. On the *Users* page, find the required user.
3. In the user's row, select the icon in the *Roles* column. A list of *Available Roles* will appear.
4. Select one or more roles.
5. Select *OK*.

> **Note**
>
> If you assign a user to a scoped role, be aware that the user is automatically assigned to all the spaces included in the scoped role. You can change the user assignment in the scoped role. See Create a Scoped Role to Assign Privileges to Users in Spaces [page 57].

Assign a Global Role to Multiple Users

> **Note**
>
> This is not relevant for scoped roles. For information about how to assign users to spaces in a scoped role, see Create a Scoped Role to Assign Privileges to Users in Spaces [page 57].

1. In the side navigation area, click (*Security*) > (*Roles*).
2. Find the role that you want to assign.
3. At the bottom of the role box, click the link *Add Users*.
4. Select one or more users from the *Assign Role to User* dialog.
5. Select *OK*.

3.3.5 Assign Users to a Role Using SAML Attributes

You can create a SAML role mapping to automatically assign users to a specific role based on their SAML attributes.

For example, you want to give a specific role to all employees that are assigned to a specific cost center. Once you've done the role mapping, if new users are assigned to the cost center in the SAML identity provider (IdP), the users will be automatically assigned to the role when logging onto SAP Datasphere via SAML authentication.

Prerequisites

Your custom SAML Identity Provider (IdP) must be configured and the authentication method selected must be *SAML Single Sign-On (SSO)* in (*System*) → (*Administration*) →*Security*. See Enabling a Custom SAML Identity Provider [page 36].

Procedure

1. In the side navigation area, click (*Security*) (*Roles*).
2. Select a role (or open the role) and click (*Open 'SAML Role Mapping'*).
3. Under *Conditions*, select a *SAML Attribute*, select a *Condition*, and enter a *Value* if required.
4. (Optional) Select **+** (*New mapping definition*) to add additional mappings to the role assignment.
   For each additional mapping, under *Conditions*, select a *SAML Attribute*, select a *Condition*, and enter a *Value* if required.
   Under *Conditions Logic*, select *AND* or *OR*.
   If *AND* is selected, the conditions for all attributes must be met for the mapping to be applied. If *OR* is selected, the conditions for only one of the attributes must be met for the mapping to be applied.
   The selected role will be applied to all users who meet the specified conditions when logging onto SAP Datasphere via SAML authentication. If the selected role was previously assigned to a user, but the user does not meet the specified conditions, the role will be revoked when the user logs in.

   > **Note**
   >
   > If a user is assigned to a scoped role via SAML attributes, the user is automatically assigned to all the spaces included in the scoped role.

   > **Caution**
   >
   > Scoped roles and all related features will be rolled out to all tenants over the course of a number of versions. For more details, see SAP Note 3380409.

   In the *Roles* page, a dedicated icon in the role tile is displayed, indicating that the users are assigned to the role via SAML attributes. When you hover over the icon, the conditions defined for the role are displayed.

3.3.6 View Authorizations (Users, Roles and Spaces)

See all the users, roles, and spaces in the tenant and how they relate to each other.

In (*Security*) (*Authorization Overview*), a user with the DW Administrator global role can see all the users, roles, and spaces in the tenant and how they relate to each other. You can filter by user, role, or space to see:

- which users are assigned with which roles to which spaces,
- which users are assigned to which global roles.

Enter a String to Search On

To display information related to the one or more terms, enter one or more characters in the *Search* field and press *Enter* (or click *Search*).

As you type, the field will begin proposing objects and search strings. Click on a string to trigger a search on it.

For example, to display all roles that are assigned to the user Lucia, enter "Lucia" in the *Search*.

Filter by Criteria

You can filter the list by any of the categories listed in the *Filter By* area of the left panel: user (in *User Name*), space (in *Scope Name*) and role (in *Role Name*).

You can select one or more values in each filter category in the *Filter By* section:

- Each value selected in a category acts as an `OR` condition. For example, to display all roles that are assigned to the users Lucia and Ahmed, select Lucia and Ahmed in the *User Name* category.
- Values selected in separate categories act together as `AND` conditions. For example, to display all the scoped roles that enables Lucia to access the Sales Asia space, select Lucia in the *User Name* category and Sales Asia in the *Scope Name* category.

3.3.7 Privileges and Permissions

A privilege represents a task or an area in SAP Datasphere and can be assigned to a specific role. The actions that can be performed in the area are determined by the permissions assigned to a privilege.

This topic contains the following sections:

Overview

A role represents the main tasks that a user performs in SAP Datasphere. Each role has a set of privileges with appropriate levels of permissions. The privileges represent areas of the application like the *Space Management* or the *Business Builder* and the files or objects created in those areas.

The standard roles provide sets of privileges and permissions that are appropriate for that role. For example, the *DW Administrator* role has all the *Spaces* permissions, while the *DW Viewer* role has none.

You can use the standard roles (see Standard Roles Delivered with SAP Datasphere [page 54]) and create your own custom roles to group together other sets of privileges and permissions (see Create a Custom Role [page 56]).

Global versus scoped privileges - Global privileges are privileges that are used at the tenant level and are not space-related, and can therefore be included in a global role, typically a tenant administrator role. Scoped privileges are privileges that are space-related and can therefore be included in a scoped role.

> ⚠ Caution
>
> Scoped roles and all related features will be rolled out to all tenants over the course of a number of versions. For more details, see SAP Note 3380409.

Global Privileges and Permissions

The following table lists the privileges and their permissions that can be included in a global role.

Global Privileges and Permissions

(`C`=Create, `R`=Read, `U`=Update, `D`=Delete, `E`=Execute, `M`=Maintain, `S`=Share, `M`=Manage)

Privilege	Permissions	Description
Spaces	C------M	Allows access to spaces in the *Space Management* tool. • *Create* - To create spaces. To perform actions on spaces, you need a combination of permissions for the privilege *Spaces* and for other privileges. See Roles and Privileges by App and Feature [page 77]. • *Manage* - To read, update and delete all spaces. **⚠ Caution** The permission *Manage* should be granted only to tenant administrators. and for other privileges. **ⓘ Note** The permissions *Read*, *Update* and *Delete* are scoped permissions and are described in the scoped privileges and permissions table (see Scoped Privileges and Permissions [page 71]). See Managing Your Space
Space Files	-------M	Allows access to all objects inside a space, such as views and tables. *Manage* - To view objects and data in all spaces. **ⓘ Note** To perform actions on spaces, you need a combination of permissions for the privilege *Spaces* and for other privileges. See Roles and Privileges by App and Feature [page 77]. **⚠ Caution** The permission *Manage* should be granted only to tenant administrators. **ⓘ Note** The permissions *Create*, *Read*, *Update* and *Delete* are scoped permissions and are described in the scoped privileges and permissions table (see Scoped Privileges and Permissions [page 71]). See Managing Your Space
Data Warehouse General	-R------	Allows users to log into SAP Datasphere. Included in all standard roles except for *DW Consumer*.

Privilege	Permissions	Description
Data Warehouse Runtime	`-R--E---`	• *Read* - Allows users of the View Analyzer to download the generated SQL analyzer plan file. See Exploring Views with View Analyzer • *Execute* - not in use
Other Datasources	`----E---`	Some connection types require this privilege. For more information, see Permissions in the *SAP Analytics Cloud Help*.
Role	`CRUD----`	Allows access to *Security* > *Roles* and *Security* > *Authorization Overview*. See Managing Roles and Privileges [page 52] and View Authorizations (Users, Roles and Spaces) [page 65]
User	`CRUD---M`	Allows access to lists of users. • `R` (*Read*) - To see a list of users in a dialog, for example when choosing which users to share a story with, or when choosing users to add to a team. • To open the *Security* > *Users* tools, you need all 4 permissions `CRUD----`. If you have only the Read permission, you cannot see the list of users in *Security* > *Users*. > **Note** > The permissions `CRUD----` are included in the DW Administrator role. When you create a custom role based on the DW Administrator role, the permissions are automatically included and you cannot edit them. • `M` (*Manage*) - To permit assigning users to roles, and approving role assignment requests from users. See Managing SAP Datasphere Users [page 45]
Activity Log	`-R-D----`	Allows access to the *Activities* page in the *Security* tool. • *Read* - To view the activities in the *Activities* page and download the activity log for a specific time period. • *Delete* - To delete the activity log for a specific time period. See Monitor Object Changes with Activities [page 207]

Privilege	Permissions	Description
Lifecycle	`-R---MS-`	Allows to import and export content via the *Content Network* and *Transport* tools. • `M` (*Maintain*) - Allows access to the tools *Content Network* and *Transport* > *Import* so the user can import packages from the *Content Network*. • `M` (*Maintain*) and `S` (*Share*) - Allows access to the tool *Transport* > *Export* so the user can export packages from *Transport* > *Export* in the *Content Network*. **Note** The permissions `-R---MS-` are included in the DW Administrator role. When you create a custom role based on the DW Administrator role, the permissions are automatically included and you cannot edit them. See Importing SAP and Partner Business Content from the Content Network and Transporting Content Between Tenants
System Information	`-RU-----`	• *Read*: To access the *About* area in the *System* tool. • *Update*: To access the *Administration*, *Configuration* and *About* areas in the *System* tool.

Catalog Asset	`CRUD---M`	• *Create*: Add an asset.
		• *Read*:
		  • Access the catalog and view asset details.
		  • Search assets, favorites, and recent.
		  • Filter linked tags, terms, and KPIs.
		• *Update*: Edit the asset name and description.
		• *Delete*: Remove an asset from the catalog.
		• *Manage*:
		  • View published and unpublished assets.
		  • Publish or unpublish assets.

Catalog Glossary	`CRUD----`	• *Create*: Use with the *Update* privilege to create a glossary.
		• *Read*: Use with *Catalog Glossary Object* to:
		  • View the term details and glossary list.
		  • Create a category.
		  • Search for terms, favorites, and assets.
		• *Update*: Edit a glossary.
		• *Delete*: Delete a glossary.

Privilege	Permissions	Description

Catalog Glossary Object	CRUD---M	• *Create*: Create a term.
		• *Read*: Use with *Catalog Glossary* to:
		  • View the term details and glossary list.
		  • Create a category.
		  • Search for terms, favorites, and assets.
		• *Update*: Edit a term.
		• *Delete*: Delete a term.
		• *Manage*:
		  • View published and unpublished terms.
		  • Publish or unpublish a term.

Catalog Tag Hierarchy	CRUD----	• *Create*: Create a tag hierarchy. • *Read*: View tag hierarchies and search for tags. • *Update*: Edit tag hierarchies. • *Delete*: Delete a tag hierarchy.
Catalog System	CRUDE---	• *Create*: Create a system. • *Read*: View systems overview. • *Update*: Configure and update a system. • *Delete*: Delete a system. • *Execute*: Synchronize source systems manually.

Catalog KPI Object	CRUD---M	• *Create*: Use with *Catalog KPI Template* with *Read* permission to create a KPI.
		• *Read*: Use with *Catalog KPI Template* with *Read* permission to:
		  • View KPI details.
		  • Search for KPIs, favorites, and recent.
		  • Filter KPIs on linked terms.
		• *Update*: Use with *Catalog KPI Template* with *Read* permission to update a KPI.
		• *Delete*: Delete a KPI.
		• *Manage*:
		  • View published and unpublished KPIs.
		  • Publish or unpublish KPIs.

Catalog KPI Template	-RU-----	• *Read*: Use with *Catalog KPI Object* with *Read* permission to:
		  • View KPI details.
		  • Search for KPIs, favorites, and recent.
		  • Filter KPIs on linked terms.
		• *Update*: Edit the KPI template.

Catalog Log	-R------	• *Read*: View and search extraction logs for assets and batch job details.

Scoped Privileges and Permissions

The following table lists the privileges and their permissions that can be included in a scoped role.

> **Note**
>
> Some permissions require others and may automatically set them. For example, setting the *Delete* permission for the *Data Warehouse Data Builder* privilege automatically sets the *Read* permission as well.

Scoped Privileges and Permissions
(C=Create, R=Read, U=Update, D=Delete, E=Execute, M=Maintain, S=Share, M=Manage)

Privilege	Permissions	Description
Spaces	-RUD----	Allows access to spaces in the *Space Management* tool. • *Read* - To view the *Space Management*. • *Update*, *Delete* - To update or delete spaces. **Note** The permissions *Create* and *Manage* are global permissions and are described in the global privileges and permissions table (see Global Privileges and Permissions [page 66]). See Managing Your Space

Privilege	Permissions	Description
Space Files	CRUD - - - -	Allows access to all objects inside a space, such as views and tables. • *Read* - To view the objects in spaces. • *Create*, *Update*, *Delete* - To create, update or delete objects in spaces. To view certain space properties or perform actions on spaces, you need a combination of permissions for the privilege *Spaces* and for other privileges. See Roles and Privileges by App and Feature [page 77]. **Note** The permission *Manage* is a global permission and is described in the global privileges and permissions table (see Global Privileges and Permissions [page 66]). See Managing Your Space
Data Warehouse Data Builder	CRUD - - S -	Allows access to all objects in the *Data Builder* app. Users with the *Share* permission can share objects to other spaces. Also allows access to the *Data Sharing Cockpit* app with the *Create*, *Read* and *Update* permissions. See Acquiring Data in the Data Builder, Preparing Data in the Data Builder and Modeling Data in the Data Builder

Privilege	Permissions	Description
Data Warehouse Remote Connection	CRUD- - - -	Allows access to remote and run-time objects: • *Read* - To view remote tables in the *Data Builder*. • *Create*, *Read*, *Update* and *Delete* - To create, update, or delete a connection in the *Connections* app, in addition to the corresponding *Space Files* permission. See Integrating Data via Connections and Acquiring Data in the Data Builder

Data Warehouse Data Integration	-RU-E- - -	Allows access to the *Data Integration Monitor* app:
		• *Read* - To view the *Data Integration Monitor*.
		• *Update*:
		  • To perform any one-off data replication/persistence actions in the *Data Integration Monitor* or *Data Builder*.
		  • To redeploy views in the *Data Builder* where data persistence is used (including in the view lineage)
		• *Execute* - To work with schedules.
		Note
		In addition to these permissions, the following *Data Integration Monitor* actions require the *Data Warehouse Data Builder* (*Read*) privilege:
		• To set up or change partitioned data loading in the *Remote Tables* monitor or in the *Views* monitor.
		• To start the *View Analyzer* in the *Views* monitor.
		See Managing and Monitoring Data Integration

Data Warehouse Business Catalog	Not in use	Not in use

Privilege	Permissions	Description
Data Warehouse Data Access Control	CRUD- - - -	Allows access to the *Data Access Control* app and objects. • *Update* - To see the *Data Access Control* app. • *Read* - To use a *Data Access Control* object to protect a view. • *Update* - To change a view, in addition to *Read*. See Securing Data with Data Access Controls
Data Warehouse Business Builder	-R- - - - - -	Allows access to the *Business Builder* app. See Modeling Data in the Business Builder
Data Warehouse Business Entity	CRUD- - - -	Allows access to business objects (dimensions and facts) defined in the *Business Builder*. See Creating a Business Entity
Data Warehouse Authorization Scenario	CRUD- - - -	Allows access to authorization scenarios defined in the *Business Builder*. Authorization scenarios are modeling abstractions for *Data Access Controls*. See Authorization Scenario
Data Warehouse Fact Model	CRUD- - - -	Allows access to fact models defined in the *Business Builder*. Fact models are shaped like consumption models but offer re-useability in other consumption models. See Creating a Fact Model
Data Warehouse Consumption Model	CRUD- - - -	Allows access to consumption models inside the *Business Builder*. Consumption models comprise perspectives which are presented as `DWC_CUBE` objects in the file repository. See Creating a Consumption Model

Privilege	Permissions	Description
Data Warehouse Folder	CRUD- - - -	Allows access to folders defined in the *Business Builder*. Folders are used to organize objects inside the *Business Builder*. See Business Builder Start Page
Data Warehouse Consumption	-RU-E- - -	Allows access to data in modeling objects: • R (*Read*) - Read data output by *Data Builder* views that have the *Expose for Consumption* switch enabled and data in *Business Builder* fact models and consumption models. Users with this permission may not preview data in local or remote tables, in views that are not exposed for consumption, in sources or intermediate nodes of graphical views (even if those views are exposed for consumption), or in *Business Builder* business entities. This permission is given to users with the standard *DW Viewer* role (who have read-only access to SAP Datasphere) and users with the *DW Consumer* role (who do not have access to SAP Datasphere and merely consume exposed data in SAP Analytics Cloud and other analytics clients). • U (*Update*) - Upload data from a csv file to a local table. For local tables with delta capture enabled, updates are tracked in the "Change Type" column. • E (*Execute*) - Access data in all *Data Builder* and *Business Builder* objects and edit data in *Data Builder* local tables. See Consuming Data Exposed by SAP Datasphere
Data Warehouse General	-R- - - - - -	Allows users to log into SAP Datasphere. Included in all standard roles except for *DW Consumer*.

Privilege	Permissions	Description
Scoped Role User Assignment	`-------M`	Allows to manage user assignment in a space. `M` (*Manage*): • To see the *Users* area in the spaces assigned to the scoped role, in addition to *Spaces* `Read`. • To edit the *Users* area in the spaces assigned to the scoped role, in addition to *Spaces* `Update`. **Note** This privilege is displayed and available for selection only in a scoped role and is selected by default in the predefined scoped role DW Scoped Space Administrator. See Create a Scoped Role to Assign Privileges to Users in Spaces [page 57]
Data Warehouse Graph Modeler	Not in use	Not in use

Permissions

The following table displays the available permissions and their definitions.

Permission	Description
Create	Permits creating new objects of this item type. Users need this permission to create spaces, views or tables, upload data into a story, or upload other local files.
Read	Permits opening and viewing an item and its content.
Update	Permits editing and updating existing items. Compare this permission with the *Maintain* permission, which doesn't allow changes to the data structure. Note: some object types need the *Maintain* permission to update data. See the Maintain entry.
Delete	Permits deletion of the item.
Execute	Permits executing the item to run a process, such as schedules.
Maintain	Permits the maintenance of data values, for example adding records to a model, without allowing changes to the actual data structure. Compare this permission with the Update permission, which does allow changes to the data structure. When granted on the Lifecycle privilege, permits importing and exporting objects.
Share	Permits the sharing of the selected item type.

Permission	Description
Manage	When granted on *Spaces* and *Space Files*, permits to view all spaces and their content (including data), regardless of whether the user is assigned to the space or not. To perform actions on spaces, you need the *Manage* permission in combination with other permissions for *Spaces* and other privileges. See Roles and Privileges by App and Feature [page 77]. **⚠ Caution** This permission should be granted only to tenant administrators.

3.3.8 Roles and Privileges by App and Feature

Review the standard roles and the privileges needed to access apps, tools, and other features of SAP Datasphere.

This topic contains the following sections:

ⓘ Note

For complete lists of standard roles, privileges, and permissions, see:

- Standard Roles Delivered with SAP Datasphere [page 54]
- Privileges and Permissions [page 65]

Apps

To access an app, tool, or editor, a user must have the following standard role or a custom role containing the listed privileges:

App	Requires Privileges (Permissions)...	Granted by Role Template...
(*Home*) See The SAP Datasphere Homepage	*Data Warehouse General* (-R------)	All roles except *DW Consumer* *DW Viewer* (read-only access)

App	Requires Privileges (Permissions)...	Granted by Role Template...
(*Repository Explorer*) See Repository Explorer	*Space Files* (-R------)	All roles except *DW Consumer* *DW Viewer* (read-only access)
(*Catalog*) See Governing and Publishing Catalog Assets	• *Catalog Asset* (CRUD---M) • *Catalog Glossary* (CRUD----) • *Catalog Glossary Object* (CRUD---M) • *Catalog Tag Hierarchy* (CRUD------) • *Catalog System* (CRUDE---) • *Catalog KPI Object* (CRUD---M) • *Catalog KPI Template* (-RU-----) • *Catalog Log* (-R------)	*Catalog Administrator* *Catalog User* (read-only access for all privileges, except no access for *Catalog System* and *Catalog Log*) In addition, the following sub-tools require the *Catalog Administrator* role: • Tag Hierarchies • Monitoring
(*Data Marketplace*) See Purchasing Data from Data Marketplace	• *Spaces* (-R------) • *Space Files* (CRUD----) • *Data Warehouse Remote Connection* (CRUD--S-) • *Data Warehouse Data Integration* (-RU-E---) • *Data Warehouse Data Builder* (CRU-----)	*DW Integrator* *DW Modeler* *DW Administrator*, *DW Space Administrator* and *DW Viewer* : read-only access

(*Semantic Onboarding*)
See Semantic Onboarding.
	Data Warehouse General (-R------)
	Each section requires a specific permission:
	• *SAP Systems*:
	  • *Data Warehouse Data Builder* (CRU-----)
	  • *Data Warehouse Business Entity* (CRU-----)
	Data Warehouse Consumption Model (CRU-----)
	• *Content Network*:
	  • *Lifecycle* (-R---MS-)
	• *Data Products* - See *Data Marketplace*, above.
		DW Viewer (read-only access)
		DW Space Administrator (all sections)
		DW Modeler (SAP Systems and *Data Products*)

App	Requires Privileges (Permissions)...	Granted by Role Template...
(*Business Builder*) Start page Dimension editor Fact editor Fact model editor Consumption model editor Authorization scenario editor See Modeling Data in the Business Builder	Each page or editor requires a separate permission: • Start page: *Data Warehouse Business Builder* (-R------) • Dimension editor: *Data Warehouse Business Entity* (CRUD----) • Fact editor: *Data Warehouse Business Entity* (CRUD----) • Fact model editor: *Data Warehouse Fact Model* (CRUD----) • Consumption model editor: *Data Warehouse Consumption Model* (CRUD----) • Authorization scenario editor: *Data Warehouse Authorization Scenario* (CRUD----) The following features need additional permissions (which are included in the *DW Modeler* role): • Preview data from any object in the *Data Preview* screen - *Data Warehouse Consumption.Execute* **Note** The *DW Viewer* role includes Data Warehouse Consumption.Read, which allows these users to preview only data from Fact models and consumption models.	*DW Space Administrator* *DW Modeler* *DW Viewer* (read-only access)

App	Requires Privileges (Permissions)...	Granted by Role Template...
(*Data Builder*) Start Page Table editor Graphical view editor SQL view editor Entity-relationship model editor Data flow editor Transformation flow editor Replication flow editor Analytic model editor Intelligent lookup editor Task chain editor See Acquiring Data in the Data Builder, Preparing Data in the Data Builder and Modeling Data in the Data Builder	All pages and editors share a single permission: • *Data Warehouse Data Builder* (CRUD--S-) The following features need additional permissions (which are included in the *DW Modeler* role): • Preview data from any object in the *Data Preview* panel - *Data Warehouse Consumption.Execute* **Note** The *DW Viewer* role includes *Data Warehouse Consumption.Read*, which allows these users to preview only data output by views with the *Expose for Consumption* switch enabled. • Upload data in a local table - *Data Warehouse Consumption.Update* • Access the local table *Data Editor* screen - *Data Warehouse Data Builder.Update* • See remote objects in *Data Builder* editors - *Data Warehouse Remote Connection.Read* The following features need additional permissions (which are included in the *DW Integrator* role): • Run an intelligent lookup - *Data Warehouse Data Integration.Update* • Run a task chain - *Data Warehouse Data Integration.Update* • Delete data in a local table - *Data Warehouse Data Integration.Update*	*DW Space Administrator* *DW Modeler* *DW Viewer* (read-only access)
(*Data Access Controls*) See Securing Data with Data Access Controls	*Data Warehouse Data Access Control* (CRUD----) **Note** The *DW Modeler* role includes *Data Warehouse Data Access Control.Read*, which allows them to apply an existing data access control to a *Data Builder* view.	*DW Space Administrator*

App	Requires Privileges (Permissions)...	Granted by Role Template...
(*Data Integration Monitor*) See Managing and Monitoring Data Integration	*Data Warehouse Data Integration* (-RU-E---) **Note** *Data Warehouse Data Integration.Update* allows you to do only manual integration tasks. The *DW Integrator* role includes *Data Warehouse Data Integration.Execute*, which also allows scheduling automated integration tasks. The following features need additional permissions (which are included in the *DW Space Administrator* role): • *Views* (monitor) > Define partitions - *Data Builder.READ* • *Views* (monitor) > *View Analyzer* - *Data Builder.READ* • *Views* (monitor) > *Generate SQL Analyzer Plan File* - *Data Warehouse.RUNTIME*	*DW Space Administrator* *DW Integrator* *DW Modeler* (manual tasks only) *DW Viewer* (read-only access)
(*Connections*) See Integrating Data via Connections	*Data Warehouse Remote Connection* (CRUD--S-) The following feature needs an additional permission (which is included in the *DW Administrator* role): • Select a location ID - *Connection.Read*	*DW Space Administrator* *DW Integrator* *DW Modeler* (read-only access) *DW Viewer* (read-only access)

Administration Tools

To access an administration tool, a user must have the following standard role or a custom role containing the listed privileges:

Tool	Requires Privileges (Permissions)...	Granted by Role Template...
(*Space Management*) See Preparing Your Space and Integrating Data	*Spaces* (CRUD---M) **Note** For detailed information on permissions for *Spaces*, see Space Management Privileges and Permissions [page 82]	*DW Administrator* (can create spaces) *DW Space Administrator* *DW Integrator* and *DW Modeler*: have read-only access to the page for their space (though they cannot see all its properties).

Tool	Requires Privileges (Permissions)...	Granted by Role Template...
(*System Monitor*) See Monitoring SAP Datasphere [page 188]	*System Information* (-RU-----)	*DW Administrator*
(*Security*) Users (see Managing SAP Datasphere Users [page 45]) Roles (see Managing Roles and Privileges [page 52]) Authorization Overview (see View Authorizations (Users, Roles and Spaces) [page 65]) Activities (see Monitor Object Changes with Activities [page 207])	The sub-tools require the following permissions: • *Users*: *User* (CRUD---M) • *Roles*: *Role* (CRUD----) • *Authorization Overview*: *Role* (CRUD----) • *Activities*: *Activity Log* (-R-D----)	*DW Administrator* (read-only access for the sub-tool *Activities*)
(*Transport*) See Transporting Content Between Tenants	*Lifecycle* (-R---MS-)	*DW Administrator* *DW Space Administrator*
(*Data Sharing Cockpit*) See Data Marketplace - Data Provider's Guide	*Data Warehouse Data Builder* (CRU-----) **Note** To create a new data provider profile, or edit an existing one, you must have the **Spaces (Update)** privilege assigned to your role. See Maintaining your Data Provider Profile.	*DW Modeler* *DW Space Administrator*
(*System*) Configuration Administration About See Administering SAP Datasphere [page 5]	*System Information* (-RU-----) **Note** • The Read (R) permission gives access to the *About* dialog. • The Update (U) permission gives access to all areas.	*DW Administrator* **Note** Users with any role can view the *About* dialog.

Space Management Privileges and Permissions

Users with different roles have different levels of access to the *Space Management* tool:

- A *DW Consumer* cannot log into SAP Datasphere.

- A *DW Viewer* can log into SAP Datasphere, but has no *Spaces* permissions. They cannot see the *Space Management* tool.
- A *DW Modeler* and a *DW Integrator* have *Spaces* `(-R------)` permission. They have read-only access to the page for their space (though they cannot see all its properties).
- A *DW Space Administrator* has *Spaces* `(-RUD----)` permissions. They can see all the space properties, and edit those outside the *General Settings* and *Workload Management* sections.
- A *DW Administrator* has *Spaces* `(CRUD---M)` permissions. They can create spaces and edit some space properties, including modifying the storage allocated and the space priority.

Various privileges and permissions are required to see and edit different parts of the *Space Management* tool:

> ### Note
>
> The global privilege *Spaces* `(-------M)` enables users to perform the following actions in all the spaces of the tenant: read, update and delete.

Action	Requires Privilege (Permission)	Granted by Role Template...
Create a Space See Create a Space [page 99]	Global privileges *Spaces* `(C------M)` and *User* `(-R------)`.	DW Administrator

Action	Requires Privilege (Permission)	Granted by Role Template...
View Space Properties	Global privilege *Spaces* (- - - - - - - M) or scoped privilege *Spaces* (- R - - - - - -) **Note** In addition, you also need the following permissions to view these properties: • *Users*: Global privileges *Role* (- R - - - - - -) or scoped privileges *Scoped Role User Assignment* (- - - - - - - M) • *Data Consumption* and *Database Users*: Global privilege *Spaces* (- - - - - - - M) or scoped privilege *Spaces* (- R - - - - - -) • *HDI Containers*:Scoped privileges *Spaces* (- R - - - - - -) and *Remote Connection* (- R - - - - - -) **Note** A DW Administrator cannot see the *HDI Containers* area in a space. • *Time Data*: Scoped privileges *Spaces* (- R - - - - - -) and *Data Builder* (- R - - - - - -) **Note** A DW Administrator cannot see the *Time Data* area in a space. • *Auditing*: Global privilege *Spaces* (- - - - - - - M) or scoped privilege *Spaces* (- - R - - - - -)	DW Administrator and DW Space Administrator **Note** A user with the role DW Modeler or DW Integrator have read-only access to the page for their space but they cannot view all its properties.
Modify *General Settings* (except for *Storage Assignment*) See Create a Space [page 99]	Global privilege *Spaces* (- - - - - - - M) or scoped privilege *Spaces* (- R U - - - - -)	DW Administrator and DW Space Administrator
Modify *Storage Assignment*, *Data Lake Access*, *Workload Management* See Create a Space [page 99], Allocate Storage to a Space [page 101] and Set a Priority and Statement Limits for a Space [page 103]	Global privilege *Spaces* (- - - - - - - M)	DW Administrator

Action	Requires Privilege (Permission)	Granted by Role Template...
Modify *Users*	Global privileges *Spaces* (- - - - - - - M) and *Role* (- - - - - - - M) or scoped privileges *Spaces* (- - U - - - - -) and *Scoped Role User Assignment* (- - - - - - - M)	DW Administrator and DW Space Administrator
Modify *Data Consumption* and *Database Users* See Create a Database User	Global privilege *Spaces* (- - - - - - - M) or scoped privileges *Spaces* (- R U - - - - -)	DW Administrator, DW Space Administrator **Note** A user with the DW Integrator role needs in addition the privilege *Spaces* (- - U - - - - -) to create database users.
Modify *HDI Containers* See Prepare Your HDI Project for Exchanging Data with Your Space	Scoped privileges *Spaces* (- - U - - - - -) and *Remote Connection* (- - U - - - - -)	DW Space Administrator **Note** A DW Administrator cannot access the *HDI Containers* area in a space.
Modify *Time Data* See Create Time Data and Dimensions	To update time data: scoped privileges *Spaces* (- - U - - - - -) and *Data Builder* (- - U - - - - -) To delete time data: scoped privileges *Spaces* (- - U - - - - -) and *Data Builder* (- - - D - - - -)	DW Space Administrator **Note** A DW Administrator cannot access the *Time Data* area in a space.
Modify *Auditing* See Enable Audit Logging	Global privilege *Spaces* (- - - - - - - M) or scoped privilege *Spaces* (- R U - - - - -)	DW Administrator and DW Space Administrator
Monitor a Space See Monitor Your Space Storage Consumption	Scoped privilege *Spaces* (- R - - - - - -) **Note** With *Spaces* (- R - - - - - M), a user with the DW Administrator role can monitor the tenant and space storage by seeing the bars about used/assigned disk/memory at the top of the *Space Management* tool (see Monitor Tenant and Space Storage [page 106]).	DW Administrator, DW Space Administrator, DW Integrator and DW Modeler

Action	Requires Privilege (Permission)	Granted by Role Template...
Lock or Unlock a Space See Lock or Unlock Your Space	Global privileges *Spaces* (-------M) or scoped privilege *Spaces* (--U-----)	DW Administrator and DW Space Administrator
Delete a Space See Delete Your Space	Global privileges *Spaces* (-------M) and *User* (-------M) or scoped privileges *Spaces* (---D----) and *Scoped Role User Assignement* (-------M)	DW Administrator and DW Space Administrator **Note** A user with a space administrator role can delete only the spaces they're assigned via a scoped role. A user with a tenant administrator role can delete any space as *Spaces* (-------M) is included in the role.

Catalog Role Privilege Dependencies

When creating a custom role for using or administering the catalog, you must set the permissions for the privileges in certain ways so that you can complete various tasks. Review the following table of tasks to see which permissions and privilege combinations you need.

To be able to access the *Catalog* app from the side navigation, all custom catalog roles need the *Read* permission on Catalog Asset.

Note

All custom catalog roles need the SAP Datasphere read permission on Space Files to allow users to mark assets, terms, and KPIs as their favorite.

Category	What do you want to do	Required combination of privileges
Assets	Search for an asset and view the detailed information for it. See Finding and Accessing Data in the Catalog	Catalog Asset: (-R------)

Category	What do you want to do	Required combination of privileges
Assets	View detailed information for an asset, including the details for any term, tag, or KPI that is linked. See Evaluating Catalog Assets	Catalog Asset: (-R------) Catalog Glossary: (-R------) Catalog Glossary Object: (-R------) Tag Hierarchy: (-R------) Catalog KPI Object: (-R------) Catalog KPI Template: (-R------)
Assets	Edit the name of the asset that appears in the catalog. See Editing and Enriching Catalog Assets	Catalog Asset: (-RU-----) Catalog Tag Hierarchy: (-R------)
Assets	Add a catalog description for the asset. See Editing and Enriching Catalog Assets	Catalog Asset: Read, Update Catalog Tag Hierarchy: (-R------)
Assets	Add a term, tag, or KPI relationship to the asset from the asset's detailed information page. See Editing and Enriching Catalog Assets	Catalog Asset: Read, Update Catalog Tag Hierarchy: (-R------) Catalog Glossary Object: (-R------) Catalog KPI Object: (-R------)
Assets	Create or delete a relationship between a tag and an asset. See Add and Manage Asset Tags	Catalog Asset: (-RU-----) Catalog Tag Hierarchy: (-R------)
Assets	Manage the relationship for a term and an asset. See Create and Manage Glossary Terms	Catalog Glossary Object: (-R------) Catalog Asset: (-RU-----)
Assets	Manage the relationship for a KPI and an asset. See Create and Manage Key Performance Indicators	Catalog KPI Object: (-R------) Catalog Asset: (-RU-----)
Assets	Publish/Unpublish assets to the catalog. See Publishing to the Catalog	Catalog Asset: (-R-----M)
Tags	Add a term, tag, or KPI relationship to the asset from the asset's detailed information page. See Editing and Enriching Catalog Assets	Catalog Asset: (-RU-----) Catalog Tag Hierarchy: (-R------) Catalog Glossary Object: (-R------) Catalog KPI Object: (-R------)
Tags	Create a tag hierarchy. See Manage Hierarchies and Tags	Catalog Tag Hierarchy: (CRU-----)

Category	What do you want to do	Required combination of privileges
Tags	Edit a tag hierarchy. See Manage Hierarchies and Tags	Catalog Tag Hierarchy: (-RU-----)
Tags	Delete a tag hierarchy. See Manage Hierarchies and Tags	Catalog Tag Hierarchy: (-R-D----)
Tags	Create or delete a relationship between a tag and an asset. See Add and Manage Asset Tags	Catalog Asset: (-RU-----) Catalog Tag Hierarchy: (-R------)
Glossary	Create a glossary. See Create and Manage a Glossary	Catalog Glossary: (C-------)
Glossary	Edit a glossary. See Create and Manage a Glossary	Catalog Glossary: (-RU-----)
Glossary	Delete a glossary. See Create and Manage a Glossary	Catalog Glossary: (-R-D----)
Glossary	Create a glossary category. See Create and Manage a Glossary Category	Catalog Glossary: (-R------) Catalog Glossary Object: (C-------)
Glossary	Edit a glossary category. See Create and Manage a Glossary Category	Catalog Glossary: (-R------) Catalog Glossary Object: (-RU-----)
Glossary	Delete a glossary category. See Create and Manage a Glossary Category	Catalog Glossary: (-R------) Catalog Glossary Object: (-R-D----)
Terms	Create a glossary term. See Create and Manage Glossary Terms	Catalog Glossary: (-R------) Catalog Glossary Object: (C-------)
Terms	Edit a glossary term. See Create and Manage Glossary Terms	Catalog Glossary: (-R------) Catalog Glossary Object: (-RU-----)
Terms	Delete a glossary term. See Create and Manage Glossary Terms	Catalog Glossary: (-R------) Catalog Glossary Object: (-R-D----)
Terms	Publish or unpublish a glossary term. See Create and Manage Glossary Terms	Catalog Glossary Object: (-R-----M)
Terms	Manage the relationship for a term and an asset. See Create and Manage Glossary Terms	Catalog Glossary Object: (-R------) Catalog Asset: (-RU-----)

Category	What do you want to do	Required combination of privileges
KPIs	Create a KPI. Create and Manage Key Performance Indicators	Catalog KPI Object: (C-------) Catalog KPI Template: (-R------)
KPIs	Edit a KPI. Create and Manage Key Performance Indicators	Catalog KPI Object: (-RU-----) Catalog KPI Template: (-R------)
KPIs	Delete a KPI. Create and Manage Key Performance Indicators	Catalog KPI Object: (-R-D----)
KPIs	Publish/Unpublish a KPI. Create and Manage Key Performance Indicators	Catalog KPI Object: (-R-----M)
KPIs	Manage the relationship for a KPI and an asset. Create and Manage Key Performance Indicators	Catalog KPI Object: (-R------) Catalog Asset: (-RU-----)
KPIs	Create a KPI category. Create and Manage Key Performance Indicator Categories	Catalog KPI Object: (C-------) Catalog KPI Template: (-R------)
KPIs	Edit a KPI category. Create and Manage Key Performance Indicator Categories	Catalog KPI Object: (-RU-----) Catalog KPI Template: (-R------)
KPIs	Delete a KPI category. Create and Manage Key Performance Indicator Categories	Catalog KPI Object: (-R-D----) Catalog KPI Template: (-R------)
KPIs	Edit KPI template. SeeDefine the Key Performance Indicator Template	Catalog KPI Template: (-RU-----)

Here are a few examples for catalog roles and permissions.

For users who...	Include these privileges in the custom role
Review assets: update asset names and descriptions, add tags, and publishe assets.	Catalog Asset: (-RU----M) Catalog Tag Hierarchy: (-RU-----)
Manage and publish glossaries, terms, and KPIs; also add terms and KPI relationships to assets.	Catalog Asset: (-RU-----) Catalog Glossary: (CRUD----) Catalog Glossary Object: (CRUD---M)

For users who...	Include these privileges in the custom role
	Catalog KPI Object: (CRUD---M)
Manage terms within existing glossaries and manages tags, but do not add these relationships to assets.	Catalog Asset: (-R------) Catalog Glossary: (-R------) Catalog Glossary Object: (CRUD---M) Catalog Tag Hierarchy: (CRUD----)

External Data Consumption

Users can consume data exposed by SAP Datasphere if they are assigned to a space via a scoped role and have the Space Files.Read permission.

Action	Requires Privileges (Permissions)...	Granted by Role Template...
Consume data in SAP Analytics Cloud, Microsoft Excel, and other clients, tools, and apps See Consuming Data Exposed by SAP Datasphere	*Space Files* (-R------)	All roles **Note** If a user does not need to access SAP Datasphere itself, and only wants to consume data exposed by it, they should be granted the *DW Consumer* role.

The Command Line Interface

To use the command line interface (see Manage Spaces via the Command Line), a user must have the following standard role or a custom role containing the listed privileges:

Command	Requires Privileges (Permissions)...	Contained in Standard Role
`datasphere dbusers`	*Spaces* (-RU-----)	*DW Administrator*
`datasphere marketplace`	*Data Builder* (CRUD----)	*DW Modeler*

Command	Requires Privileges (Permissions)...	Contained in Standard Role
`datasphere objects`	• *Data Builder* (CRUD----) • *Data Warehouse Business Entity* (CRUD----) • *Data Warehouse Fact Model* (CRUD----) • *Data Warehouse Consumption Model* (CRUD----) • *Data Warehouse Authorization Scenario* (CRUD----)	*DW Modeler*
`datasphere scoped-roles`	*Role* (CRUD----)	*DW Administrator*

`datasphere spaces`	• Create a space and set storage, priority:
	  • *Spaces* `(C------M)`
	  • *User* `(-R------)`
	• Update/ delete spaces:
	  • *Spaces* (-RUD---M)
	• Update space users:
	  • *Team* (-RUD---M)
	  • *Scoped Role User Assignment* (-------M)
		DW Administrator

`datasphere tasks`	*Data Warehouse Data Integration* (-RU-E---)	*DW Integrator*
`datasphere users`	*User* (CRUD---M)	*DW Administrator*

3.3.9 Automated Conversion to Scoped Roles

For SAP Datasphere tenants that were created before version 2023.21, the roles and user assignment to spaces have been converted so that users can continue to perform the same actions as before in their spaces.

This topic contains the following sections:

> ⚠ Caution
>
> Scoped roles and all related features will be rolled out to all tenants over the course of a number of versions. For more details, see SAP Note 3380409.

Introduction to Automated Conversion to Scoped Roles

The way a DW Administrator gives privileges to users to do certain actions in spaces has changed.

Before and After Conversion

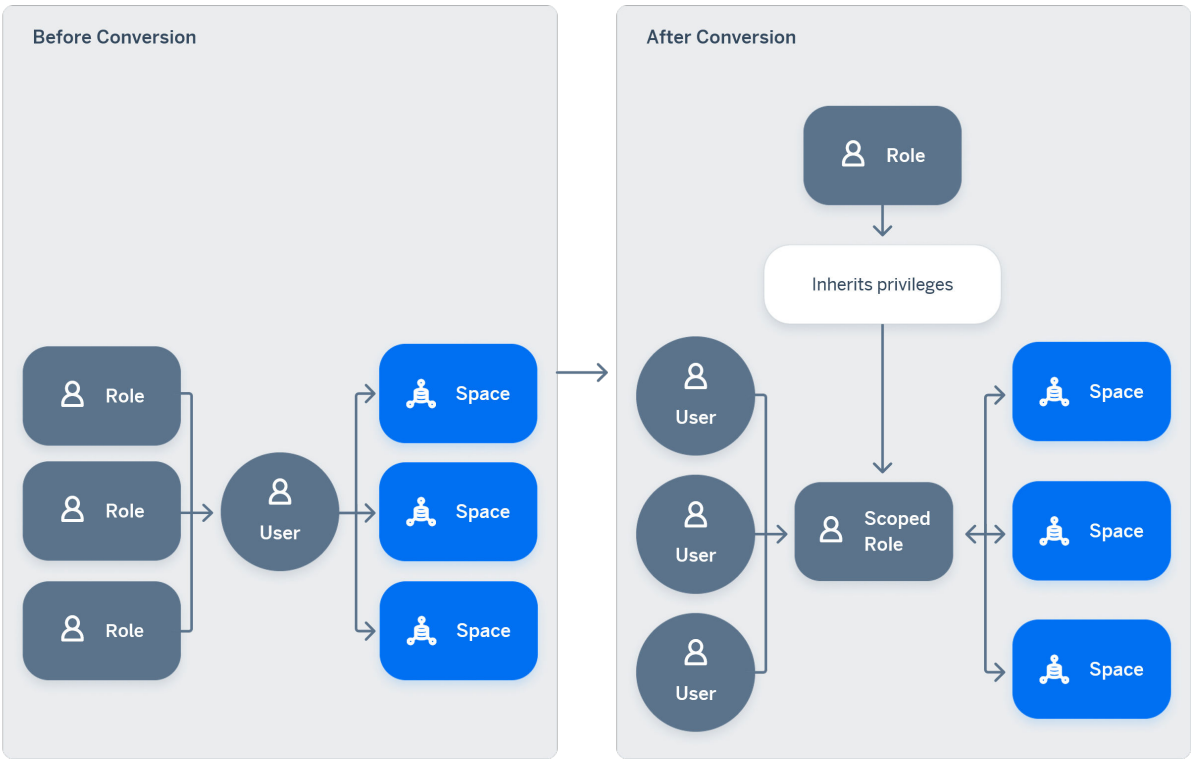

Before Conversion	After Conversion
A DW Administrator assigned a role to a user and assigned the user as a member of a space. As a consequence: • A user had the same one or more roles in all the spaces he was a member of. • A DW Administrator assigned users space by space by going in each space page.	A DW Administrator assigns a role to one or more users and one or more spaces within a new role: a scoped role. As a consequence: • A user can have different roles in different spaces: be a modeler in space Sales Germany and Sales France and a viewer in space Europe Sales. • A DW Administrator can give a role to many users in many spaces, all in one place in a scoped role. See Create a Scoped Role to Assign Privileges to Users in Spaces [page 57]. A DW Space Administrator can then manage users in their spaces and the changes are reflected in the scoped roles. See Control User Access to Your Space.

Converted Roles

You can now use global roles for tenant-wide actions and scoped roles for space-related actions.

Role Conversion

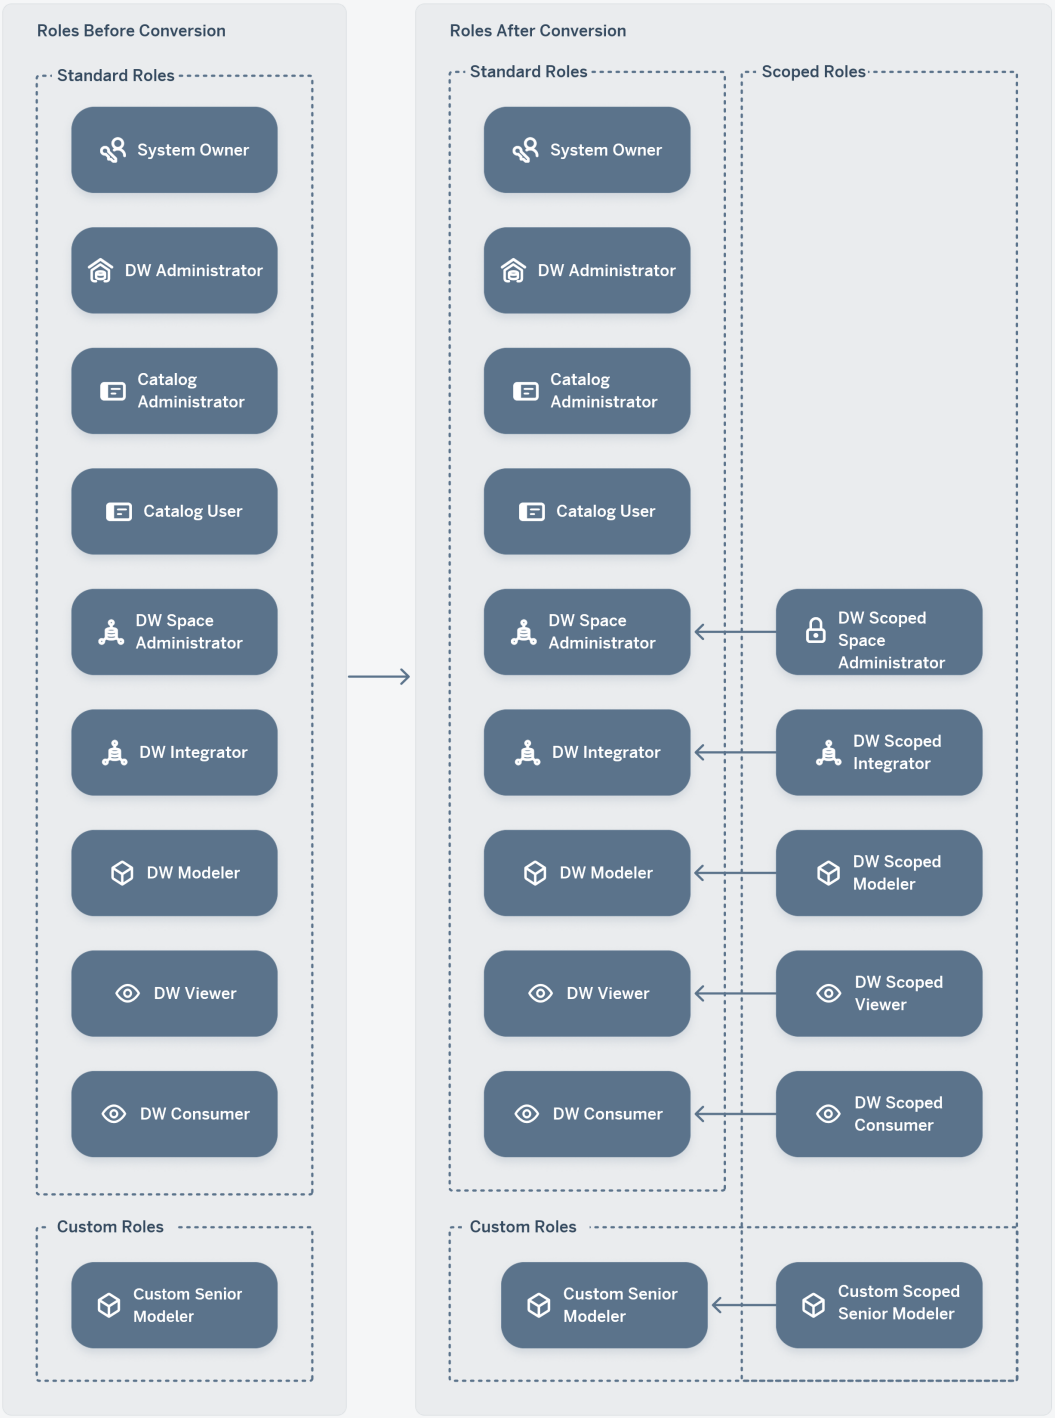

The *Roles* page lists the same standard and custom roles as before the conversion, and in addition the scoped roles that have been automatically created.

- DW Administrator, Catalog Administrator and Catalog User: these standard roles are considered as global roles. They now include only privileges that are global, which means privileges that apply to the tenant and are not space-related. For example, the DW Administrator role no more grants access to any of the modeling apps of SAP Datasphere (such as Data Builder).
  Users who previously had these roles are still assigned to them after conversion.
  Users who previously had the DW Administrator role and were members of certain spaces are assigned to the new DW Scoped Space Administrator role for those spaces they previously had access to.
  The user who previously had the System Owner role and was member of certain spaces is assigned to the new DW Scoped Space Administrator role for those spaces the user previously had access to.
- A single scoped role is created for each standard role (outside of DW Administrator, Catalog Administrator and Catalog User) and each custom role and all the users who previously had that standard or custom role are assigned to the new scoped role but only for those spaces they previously had access to.

  > ⓘ Note
  >
  > All the spaces of the tenant are included in each scoped role created, but not all users are assigned to all spaces. See the example of scoped role below.

  For each standard or custom role, two roles are available after the conversion: the initial standard or custom role (which acts as a template for the scoped role) and the scoped role created.
  Each scoped role includes privileges which are now considered as scoped privileges.
- Users who previously had the DW Space Administrator role are assigned to these 2 roles: the standard role DW Space Administrator and the new scoped role DW Scoped Space Administrator. Users who manage spaces primarily need scoped permissions to work with spaces, but they also need some global permissions (such as Lifecycle when transporting content packages). To provide such users with the full set of permissions they need, each space administrator is assigned to the scoped role DW Scoped Space Administrator to receive the necessary scoped privileges, and they are also assigned directly to the DW Space Administrator role in order to receive the additional global privileges.

> ⓘ Note
>
> - Specific case - no role assigned to a user: Before conversion, a DW Administrator assigned a user to certain spaces but did not assign a role to the user. As no role was assigned to the user, the user-to-spaces assignment is not kept after conversion.
> - Privileges and permissions are now either global or scoped. See Privileges and Permissions [page 65].

Example of a Converted Scoped Role

In this example, users assigned to a custom role called « Senior Modeler » were members of certain spaces before the conversion, as shown below.

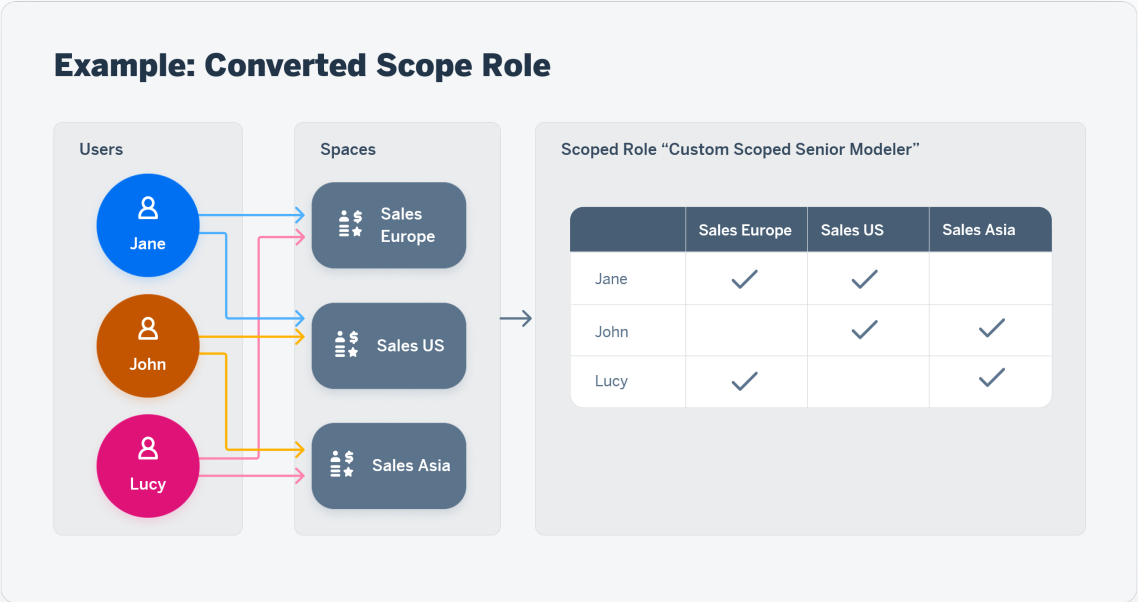

The custom role « Senior Modeler » has been converted to the scoped role « Custom Scoped Senior Modeler » and the users who previously had that custom role « Senior Modeler » are assigned to the scoped role but only for the spaces they previously had access to.

Adapting Converted Scoped Roles

The scoped roles that are automatically created during the conversion ensure that users can continue to perform the same actions as before the conversion. However, we recommend that you do not use the automatically created scoped roles and that you create your own scoped roles by logical groups as soon as possible.

In this example, the following scoped roles have been automatically created during conversion:

- DW Scoped Space Administrator
- DW Scoped Modeler
- DW Scoped Viewer
- DW Scoped Consumer

There are 4 spaces: Sales US, Sales Europe, Finance US and Finance Europe, which can be logically organized in one Sales group and one Finance group.

You should create a set of scoped roles for each logical group of spaces, add the relevant spaces and the relevant users and assign the users to the spaces in the scoped roles. The users will have access to the spaces with the appropriate privileges.

	Sales Spaces	Finance Spaces
Scoped Roles	• DW Sales Space Administrator • DW Sales Modeler • DW Sales Viewer • DW Sales Consumer	• DW Finance Space Administrator • DW Finance Modeler • DW Finance Viewer • DW Finance Consumer
Spaces	• Sales US • Sales Europe	• Finance US • Finance Europe

For more information about creating a scoped role, see Create a Scoped Role to Assign Privileges to Users in Spaces [page 57].

> **Note**
>
> In addition to the standard workflows, you can also create scoped roles and assign scopes and users to them via the command line (see Manage Scoped Roles via the Command Line).

3.3.10 Transfer the System Owner Role

The individual who purchases SAP Datasphere is automatically designated as the system owner. If you, as the purchaser, are not the right person to administer the system, you can transfer the system owner role to the appropriate person in your organization.

Prerequisites

You must be logged on as a user with the *System Information Update* privilege.

> **Note**
>
> Transferring the system owner role is **not** possible if you only have one license for SAP Datasphere.

Context

1. On the *Users* page of the *Security* area, select the user you want to assign the system owner role to.
2. Select *(Assign as System Owner)*.
   The *Transfer System Owner Role* dialog appears.
3. Under *New Role*, enter a new role for the previous system owner, or select to open a list of available roles.

   > **Note**
   >
   > One or more roles may be selected.

4. Select *OK*.

4 Creating Spaces and Allocating Storage

All data acquisition, preparation, and modeling in SAP Datasphere happens inside spaces. A space is a secure area - space data cannot be accessed outside the space unless it is shared to another space or exposed for consumption.

An administrator must create one or more spaces. They allocate disk and memory storage to the space, set its priority, and can limit how much memory and how many threads its statements can consume.

If the administrator assigns one or more space administrators via a scoped role, they can then manage users, create connections to source systems, secure data with data access controls, and manage other aspects of the space (see Managing Your Space).

4.1 Create a Space

Create a space, allocate storage, and set the space priority and statement limits.

Context

> **Note**
>
> Only administrators can create spaces, allocate storage, and set the space priority and statement limits. The remaining space properties can be managed by the space administrators that the administrator assigns to the space via a scoped role.

Procedure

1. In the side navigation area, click  (*Space Management*), and click *Create*.
2. In the *Create Space* dialog, enter the following properties, and then click *Create*:

Property	Description
Space Name	Enter the business name of the space. Can contain a maximum of 30 characters, and can contain spaces and special characters.

Property	Description
Space ID	Enter the technical name of the space. Can contain a maximum of 20 uppercase letters or numbers and must not contain spaces or special characters other than _ (underscore). Unless advised to do so, must not contain prefix _SYS and should not contain prefixes: DWC_, SAP_ (See Rules for Technical Names [page 106]).

The space is created and its property sheet opens.

3. In the *General Settings* section, review the following properties:

Property	Description
Space ID	Enter the technical name of the space. Can contain a maximum of 20 uppercase letters or numbers and must not contain spaces or special characters other than _ (underscore). Unless advised to do so, must not contain prefix _SYS and should not contain prefixes: DWC_, SAP_ (See Rules for Technical Names [page 106]).
Space Name	Enter the business name of the space. Can contain a maximum of 30 characters, and can contain spaces and special characters.
Space Status	[read-only] Displays the status of the space. Newly-created spaces are always active.
Space Type	[read-only] Displays the type of the space. You can only create spaces of type SAP Datasphere.
Created By	[read-only] Displays the user that created the space.
Created On	[read-only] Displays the date and time when the space was created.
Deployment Status	[read-only] Displays the deployment status of the space. Newly-created spaces are deployed, but when you make changes, you need to save and re-deploy them before they are available to space users.
Deployed On	[read-only] Displays the date and time when the space was last deployed.

4. [optional] Use the *Space Storage* properties to allocate disk and memory storage to the space and to choose whether it will have access to the SAP HANA data lake.

   For more information, see Allocate Storage to a Space [page 101].

5. [optional] Use the properties in the *Workload Management* section to prioritize between spaces for resource consumption and set limits to the amount of memory and threads that a space can consume when processing statements.

   For more information, see Set a Priority and Statement Limits for a Space [page 103].

6. [optional] Use the remaining sections to further configure the space.

   - *Data Access/Data Consumption*: Modify the following property, if appropriate:

Property	Description
Expose for Consumption by Default	Choose the default setting for the *Expose for Consumption* property for views created in this space.

   - *Data Access/Database Users* - Use the list in the *Database Users* section to create users who can connect external tools and read from and write to the space. See Create a Database User.
   - *Data Access/HDI Containers* - Use the list in the *HDI Containers* section to associate HDI containers to the space. See Prepare Your HDI Project for Exchanging Data with Your Space.

> **Note**
>
> A user with the DW Administrator role only cannot see the *HDI Containers* area.

- *Time Data/Time Tables and Dimensions* - Click the button in the *Time Tables and Dimensions* section to generate time data in the space. See Create Time Data and Dimensions.

  > **Note**
  >
  > A user with the DW Administrator role only cannot see the *Time Tables and Dimensions* area.

- *Auditing/Space Audit Settings* - Use the properties in the *Space Audit Settings* section to enable audit logging for the space. See Enable Audit Logging.

7. Click *Deploy* to deploy your space to the database.
8. Add your space to one or more scoped roles by doing one of the following actions:
   - Add your space to an existing scoped role (see Add Spaces to a Scoped Role [page 61]).
   - Create a scoped role and add your space and at least one user to the scoped role (see Create a Scoped Role [page 60]).

   For more information, see Create a Scoped Role to Assign Privileges to Users in Spaces [page 57]).

   All users assigned to the space via the scoped roles are automatically displayed in the *Users* area of the space page. In this area, you can add or remove users to/from scoped roles for your space (see Control User Access to Your Space. Either an administrator or a user with space administrator privileges can do so.

4.2 Allocate Storage to a Space

Use the *Space Storage* properties to allocate disk and memory storage to the space and to choose whether it will have access to the SAP HANA data lake.

Context

SAP Datasphere supports data tiering using the features of SAP HANA Cloud:

- Memory Storage (hot data) - Keep your most recent, frequently-accessed, and mission-critical data loaded constantly in memory to maximize real-time processing and analytics speeds.
  When you persist a view, the persisted data is stored in memory (see Persist View Data).
- Disk (warm data) - Store master data and less recent transactional data on disk to reduce storage costs. When you load data to a local table or replicate data to a remote table in SAP Datasphere, the data is stored on disk by default, but you can load it in memory by activating the *Store Table Data in Memory* switch (see Accelerate Table Data Access with In-Memory Storage).
- Data Lake (cold data) - Store historical data that is infrequently accessed in the data lake. With its low cost and high scalability, the data lake is also suitable for storing vast quantities of raw structured and unstructured data, including IoT data. For more information, see Integrating Data to and From SAP HANA Cloud Data Lake.

You can allocate specific amounts of memory and disk storage to a space or disable the *Enable Space Quota* option, and allow the space to consume all the storage it needs, up to the total amount available in your tenant.

Procedure

1. In the side navigation area, click  (*Space Management*), locate your space tile, and click *Edit* to open it.
2. Use the *Space Storage* properties to allocate disk and memory storage to the space and to choose whether it will have access to the SAP HANA data lake.

Property	Description
Enable Space Quota	Disable this option to allow the space to consume any amount of disk and memory storage up to the total amounts available in your tenant. If this option was disabled and then subsequently re-enabled, the *Disk* and *Memory* properties are initialized to the minimum values required by the current contents of the space. Default: Enabled
Disk (GB)	Enter the amount of disk storage allocated in GB. You can use the buttons to change the amount by whole GBs or enter fractional values in increments of 100 MB by hand. Default: 2 GB
Memory (GB)	Enter the amount of memory storage allocated in GB. This value cannot exceed the amount of disk storage allocated. You can use the buttons to change the amount by whole GBs or enter fractional values in increments of 100 MB by hand. **Note** The memory allocated is used to store data and is not related to processing memory. For more information on limiting processing memory in a space, see Set a Priority and Statement Limits for a Space [page 103]. Default: 1 GB
Use This Space to Access the Data Lake	Enable access to the SAP HANA Cloud data lake. Enabling this option is only possible if no other space already has access to the data lake. Default: Disabled

> **Note**
>
> If a space exceeds its allocations of memory or disk storage, it will be locked until a user of the space deletes the excess data or an administrator assigns additional storage. See Lock or Unlock Your Space.

3. Click *Save* to save your changes to the space, or *Deploy* to save and immediately make the changes available to users assigned to the space.

4.3 Set a Priority and Statement Limits for a Space

Use the properties in the *Workload Management* section to prioritize between spaces for resource consumption and set limits to the amount of memory and threads that a space can consume when processing statements.

Procedure

1. In the side navigation area, click  (*Space Management*), locate your space tile, and click *Edit* to open it. Then, select *Workload Management*.
2. To prioritize between spaces, enter in the *Space Priority* section the prioritization of this space when querying the database. You can enter a value from 1 (lowest priority) to 8 (highest priority). The default value is 5. In situations where spaces are competing for available threads, those with higher priorities have their statements run before those of spaces with lower priorities.
3. To manage other workload parameters, you can select either of the following in the *Space Configuration* dropdown list:

    - *Default*. The default configuration provides generous resource limits, while preventing any single space from overloading the system. The default configuration is applied by default to new spaces.
      These statement limit and admission control parameters are taken into account in the default configuration and cannot be changed:

Parameter Type	Parameter	Value
Admission Control	ADMISSION CONTROL QUEUE CPU THRESHOLD	80%
	ADMISSION CONTROL REJECT CPU THRESHOLD	99%
Statement Limits	TOTAL STATEMENT THREAD LIMIT	80%

    - *Custom*. These statement limit and admission control parameters are taken into account in the custom configuration. You can specify only the value for statements limits to set maximum total thread and memory limits that statements running concurrently in the space can consume:

      > ⚠ Caution
      >
      > Be aware that changing the statement limits may cause performance issues.

Parameter Type	Parameter	Value
Admission Control	ADMISSION CONTROL QUEUE CPU THRESHOLD	80%

Parameter Type	Parameter	Value
	ADMISSION CONTROL REJECT CPU THRESHOLD	99%
Statement Limits	TOTAL STATEMENT THREAD LIMIT	In the *Total Statement Thread Limit* area, enter the maximum number (or percentage) of threads that statements running concurrently in the space can consume. You can enter a percentage between 1% and 70% (or the equivalent number) of the total number of threads available in your tenant. Setting this limit prevents the space from consuming too many threads, and can help with balancing resource consumption between competing spaces. **⚠ Caution** Be aware that setting this limit too low may impact statement performance, while excessively high values may impact the performance of statements in other spaces. **ⓘ Note** If you've set a thread limit percentage (or number) that is no longer allowed, it is changed to the default percentage 70% (or to the number corresponding to 70% of the total number of threads). Default: 70%
	TOTAL STATEMENT MEMORY LIMIT	In the *Total Statement Memory Limit* area, enter the maximum number (or percentage) of GBs of memory that statements running concurrently in the space can consume. You can enter any value or percentage between 0 (no limit) and the total amount of memory available in your tenant. Setting this limit prevents the space from consuming all available memory, and can help with balancing resource consumption between competing spaces. **⚠ Caution** Be aware that setting this limit too low may cause out-of-memory issues, while excessively high values or 0 may allow the space to consume all available system memory. Default: 80%

ⓘ Note

Admission control is designed to avoid overloading the system under peak load by denying any further SQL requests when the load on the system is equal to or exceeds a given threshold.

You can investigate why statements are being queued or rejected.

- Events related to requests which have been queued for longer than 5 seconds are logged and can be reviewed in the M_ADMISSION_CONTROL_EVENTS view. For more information, see Managing Peak Load (Admission Control) in the *SAP HANA Cloud, SAP HANA Database Administration Guide*.
- You can monitor the statements that are queued or rejected by viewing the cards dedicated to admission control in the *Dashboard* tab of the *System Monitor*. For more information, see Monitoring SAP Datasphere [page 188].

A statement which exceeds a reject threshold is rejected with the SQL error 616: 'rejected by workload class configuration'. A statement which exceeds a queue threshold is queued for up to 10 minutes, after this time the statement is rejected with the SQL error 616: 'queue wait timeout exceeded'. For more information, see Properties for Workload Classes and Mappings in the *SAP HANA Cloud, SAP HANA Database Administration Guide*.

Note

If too many statements are rejected, we recommend you to do these two actions:

- **Decrease the total statement thread limit for the spaces which consume a large amount of CPU time**.
  First, identify the spaces which consume a large amount of CPU time: As a database analysis user, analyze the M_WORKLOAD_CLASS_STATISTICS view in the Database Explorer, like in this example:

```
SELECT "MD"."SPACE_ID", "WCS"."TOTAL_STATEMENT_CPU_TIME",
"WCS"."TOTAL_STATEMENT_REJECT_COUNT"
FROM "SYS"."M_WORKLOAD_CLASS_STATISTICS" AS "WCS"
LEFT JOIN "DWC_TENANT_OWNER"."SPACE_METADATA" AS "MD" ON
"WCS"."WORKLOAD_CLASS_NAME" = "MD"."VALUE"
AND "MD"."SECTION" = '_workloadManagement' AND "MD"."KEY" =
'workloadClassName'
WHERE "MD"."SPACE_ID" IS NOT NULL AND "MD"."SPACE_ID" != '$$global$
$' ORDER BY "WCS"."TOTAL_STATEMENT_CPU_TIME" DESC
```

  Using this sample code, all the spaces can be listed by the TOTAL_STATEMENT_CPU_TIME descending order, which enables you to identify the spaces that consumed the most CPU time. As a second step, go to the *Workload Configuration* area of each identified space, select the configuration *Custom* and decrease the total statement thread limit. Some statements will take longer to run but will not be rejected.
- **Avoid that tasks which consume a high load of CPU run at the same time**.
  You can adjust the task schedules in the *Data Integration Monitor*. See Scheduling Data Integration Tasks.

4. Click *Save* to save your changes to the space, or *Deploy* to save and immediately make the changes available to users assigned to the space.

4.4 Monitor Tenant and Space Storage

You can see the total storage available and the amount assigned to and used by spaces in the bars at the top of the *Space Management* page.

> **Note**
>
> You can also see the information below in the *System Monitor*. For more information, see Monitoring SAP Datasphere [page 188].

The following information is available:

- *Disk Used for Storage* - Shows the total amount of disk storage used. Hover over this bar to see a breakdown between:
  - *Data in Spaces*: All data that is stored in spaces.
  - *Audit Log Data*: Data related to audit logs (see Audit Logging).

    > **Note**
    >
    > Audit logs can grow quickly and consume a great deal of disk storage (see Delete Audit Logs [page 206]).

  - *Other Data*: Includes data stored in database user group schemas (see Creating a Database User Group [page 180]) and SAP HANA data (such as statistics schemas).
  - *Administrative Data*: Data used to administer the tenant and all spaces (such as space quota, space version). Includes all information stored in the central schemas (DWC_GLOBAL, DWC_GLOBAL_LOG, DWC_TENANT_OWNER).
- *Disk Assigned for Storage* - Shows the total amount of disk storage assigned to all spaces.
- *Memory Used for Storage* -Shows the total amount of memory storage used in all spaces.
- *Memory Assigned for Storage* - Shows the total amount of memory storage assigned to all spaces.

4.5 Rules for Technical Names

Rules and restrictions apply to the technical names of objects that you create in SAP Datasphere. The technical name by default is synchronized with the business name by using rules to automatically replace invalid characters.

When specifying the technical name of an object, bear in mind the following rules and restrictions:

Object Type	Rule	Maximum Length
Space	The space ID can only contain uppercase letters, numbers, and underscores (_). Reserved keywords, such as SYS, CREATE, or SYSTEM, must not be used. Unless advised to do so, the ID must not contain prefix _SYS and should not contain prefixes: DWC_, SAP_. The maximum length is 20 characters. Reserved keywords: SYS, PUBLIC, CREATE, SYSTEM, DBADMIN, MONITORING, PAL_STEM_TFIDF, SAP_PA_APL, DWC_USER_OWNER, DWC_TENANT_OWNER, DWC_AUDIT_READER, DWC_GLOBAL, and DWC_GLOBAL_LOG. Also, the keywords that are reserved for the SAP HANA database cannot be used in a space ID. See Reserved Words in the *SAP HANA SQL Reference Guide for SAP HANA Platform*.	20
SAP BW bridge instance Remote table generated during the import of analysis authorizations from a SAP BW or SAP BW/4HANA system	The technical name can contain any characters except for the asterisk (*), colon (:), and hash sign (#). Also, tab, carriage return, and newline must not be used, and space must not be used at the start of the name. The maximum length is 50 characters.	50
Entity created in the Data Builder and Business Builder, for example a table or view Column Attribute Measure Analytical measure	The technical name can only contain alphanumeric characters and underscores (_). The maximum length is 50 characters.	50
Association	The technical name can only contain alphanumeric characters, underscores (_), and dots (.). The maximum length is 10.	10
Input parameter	The technical name can only contain uppercase letters, numbers, and underscores (_). The maximum length is 50 characters.	30
Database analysis user	The user name suffix can only contain uppercase letters, numbers, and underscores (_). The maximum length is 41 characters. This suffix is added to the default prefix DWCDBUSER# to create your full user name. Note that you cannot change the prefix as it is a reserved prefix.	31 (40 minus prefix)
Database user group user	The user name suffix can only contain uppercase letters, numbers, and underscores (_). The maximum length is 41 characters. This suffix is added to the default prefix DWCDBGROUP# to create your full user name. Note that you cannot change the prefix as it is a reserved prefix.	30 (40 minus prefix)

Object Type	Rule	Maximum Length
Database user (Open SQL schema)	The user name suffix can only contain uppercase letters, numbers, and underscores (_). The maximum length is 41 characters. This suffix is added to the default prefix <space ID># to create your full user name. Note that you cannot change the prefix.	40 minus space name (or 41 minus prefix)
Connection	The technical name can only contain alphanumeric characters and underscores (_). Underscore (_) must not be used at the start or end of the name. The maximum length is 40 characters.	40
Data access control	The technical name can only contain alphanumeric characters, and underscores (_). The maximum length is 50 characters.	50

The technical name by default is synchronized with the business name. While entering the business name, invalid characters are replaced in the technical name as follows:

Rule	Example
Reserved keywords which are not allowed are removed.	" SYS" ▶ ""
Leading underscores (_) are removed.	"_NAME" ▶ "NAME"
Leading and trailing whitespaces (" ") are removed.	" NAME " ▶ "NAME"
Whitespaces (" ") within a name are replaced with underscores (_).	"NA ME" ▶ "NA_ME"
Characters with diacritical signs are replaced with their basic character.	"Namé" ▶ "Name"
Non-alphanumeric characters are removed.	"N$ME" ▶ "NME"
Dots (.) and double quotes (") are replaced with underscores (_).	"N.AM"E" ▶ "N_AM_E"
Leading dots (.) are removed.	".NAME" ▶ "NAME"

4.6 Create Spaces via the Command Line

You can use the SAP Datasphere command line interface, `datasphere`, to create, read, update, and delete spaces. You can set space properties, assign (or remove) users, create database users, create or update objects (tables, views, and data access controls), and associate HDI containers to a space.

To use `datasphere` to create spaces, you must have an SAP Datasphere user with the *DW Administrator* role or equivalent permissions (see Roles and Privileges by App and Feature [page 77]).

For more information, see Manage Spaces via the Command Line.

5 Preparing Connectivity for Connections

You need to perform some preparatory steps to be able to create and use connections in SAP Datasphere. The steps depend on the source you want to connect to and on the features you want to use with the connection.

The following overview lists the most common prerequisites per connection type and points to further information about what needs to be prepared to connect and use a connection.

Connection Type	Remote Tables: Data Provisioning Agent Required?	Remote Tables: Installation of Third-Party JDBC Library Required?	Data Flows and Replication Flows: Cloud Connector Required for On-Premise Sources?	Data Flows: Third-Party Driver Upload Required?	SAP Datasphere IP Required in Source Allowlist?	Source IP Required in SAP Datasphere IP Allowlist?	Additional Information and Prerequisites
Adverity Connections	no	no	no	no	no	yes	Prepare Connectivity to Adverity [page 136]
Amazon Athena Connections	no	no	no	no	no	no	Prepare Connectivity to Amazon Athena [page 137]
Amazon Redshift Connections	yes	yes	no	yes	yes (Outbound IP Address)	no	Prepare Connectivity to Amazon Redshift [page 138]
Amazon Simple Storage Service Connections	no	no	no	no	no	no	n/a
Apache Kafka Connections	no	no	yes	no	no	no	Prepare Connectivity to Apache Kafka [page 137]
Confluent Connections	no	no	yes	no	no	no	Prepare Connectivity to Confluent [page 138]
Cloud Data Integration Connections	yes	no	yes (for data flows)	no	no	no	Prepare Connectivity for Cloud Data Integration [page 139]

Connection Type	Remote Tables: Data Provisioning Agent Required?	Remote Tables: Installation of Third-Party JDBC Library Required?	Data Flows and Replication Flows: Cloud Connector Required for On-Premise Sources?	Data Flows: Third-Party Driver Upload Required?	SAP Datasphere IP Required in Source Allowlist?	Source IP Required in SAP Datasphere IP Allowlist?	Additional Information and Prerequisites
Generic JDBC Connections	yes (for on-premise)	yes	no	no	no	no	Prepare Connectivity for Generic JDBC [page 140]
Generic OData Connections	no	no	yes (for data flows)	no	no	no	Prepare Connectivity for Generic OData [page 140]
Generic SFTP Connections	no	no	no	no	no	no	Prepare Connectivity for Generic SFTP [page 141]
Google BigQuery Connections	no	no	no	yes	no	no	Prepare Connectivity to Google BigQuery [page 142]
Google Cloud Storage Connections	no	no	no	no	no	no	n/a
Hadoop Distributed File System Connections	no	no	no	no	no	no	n/a
Microsoft Azure Blob Storage Connections	no	no	no	no	no	no	n/a
Microsoft Azure Data Lake Store Gen1 Connections (Deprecated)	no	no	no	no	no	no	n/a
Microsoft Azure Data Lake Store Gen2 Connections	no	no	no	no	no	no	n/a

Connection Type	Remote Tables: Data Provisioning Agent Required?	Remote Tables: Installation of Third-Party JDBC Library Required?	Data Flows and Replication Flows: Cloud Connector Required for On-Premise Sources?	Data Flows: Third-Party Driver Upload Required?	SAP Datasphere IP Required in Source Allowlist?	Source IP Required in SAP Datasphere IP Allowlist?	Additional Information and Prerequisites
Microsoft Azure SQL Database Connections	yes	yes	no	no	yes (Outbound IP Address)	no	Prepare Connectivity to Microsoft Azure SQL Database [page 142]
Microsoft SQL Server Connections	yes	yes	yes (for data flows)	no (pre-bundled; no upload required)	no	no	Prepare Connectivity to Microsoft SQL Server [page 143]
Open Connectors Connections	no	no	no	no	no	no	Prepare Connectivity to SAP Open Connectors [page 144]
Oracle Connections	yes	yes	no	yes	no	no	Prepare Connectivity to Oracle [page 145]
Precog Connections	no	no	no	no	no	yes	Prepare Connectivity to Precog [page 146]
SAP ABAP Connections	yes (for on-premise)	no	yes (for on-premise: for data flows and replication flows)	no	no	no	Prepare Connectivity to SAP ABAP Systems [page 146]
SAP BW Connections	yes	no	yes (for data flows)	no	no	no	Prepare Connectivity to SAP BW [page 149]
SAP BW/4HANA Model Transfer Connections	yes (for model import - to connect to the SAP HANA database of SAP BW/4HANA)	no	yes (for model import - to make http requests to SAP BW/4HANA)	no	no	no	Preparing SAP BW/4HANA Model Transfer Connectivity [page 150]

Connection Type	Remote Tables: Data Provisioning Agent Required?	Remote Tables: Installation of Third-Party JDBC Library Required?	Data Flows and Replication Flows: Cloud Connector Required for On-Premise Sources?	Data Flows: Third-Party Driver Upload Required?	SAP Datasphere IP Required in Source Allowlist?	Source IP Required in SAP Datasphere IP Allowlist?	Additional Information and Prerequisites
SAP ECC Connections	yes	no	yes (for data flows)	no	no	no	Prepare Connectivity to SAP ECC [page 153]
SAP Fieldglass Connections	yes	no	yes (for data flows)	no	no	no	Prepare Connectivity to SAP Fieldglass [page 154]
SAP HANA Connections	yes (for on-premise)	no	yes (for on-premise: for data flows and replication flows, or when using Cloud Connector for remote tables feature)	no	no	Cloud Connector IP (for on-premise when using Cloud Connector for remote tables feature)	Prepare Connectivity to SAP HANA [page 154]
SAP HANA Cloud, Data Lake Files Connections	no	no	no	no	no	no	no
SAP HANA Cloud, Data Lake Relational Engine Connections	no	no	no	no	no	no	n/a
SAP Marketing Cloud Connections	yes	no	yes (for data flows)	no	no	no	Prepare Connectivity to SAP Marketing Cloud [page 156]
SAP SuccessFactors Connections	no	no	no	no	yes (HANA IP Address)	no	Prepare Connectivity to SAP SuccessFactors [page 156]

Connection Type	Remote Tables: Data Provisioning Agent Required?	Remote Tables: Installation of Third-Party JDBC Library Required?	Data Flows and Replication Flows: Cloud Connector Required for On-Premise Sources?	Data Flows: Third-Party Driver Upload Required?	SAP Datasphere IP Required in Source Allowlist?	Source IP Required in SAP Datasphere IP Allowlist?	Additional Information and Prerequisites
SAP S/4HANA Cloud Connections	yes	no	no	no	no	no	Prepare Connectivity to SAP S/4HANA Cloud [page 157]
SAP S/4HANA On-Premise Connections	yes (for model import)	no	yes (for data flows, replication flows, and model import)	no	no	no	Prepare Connectivity to SAP S/4HANA On-Premise [page 158]

> **Note**
>
> For information about supported versions of sources that are connected via SAP HANA smart data integration and its Data Provsioning Agent, see the SAP HANA smart data integration and all its patches Product Availability Matrix (PAM) for SAP HANA SDI 2.0.
>
> For information about necessary JDBC libraries for connecting to sources from third-party vendors, see:
>
> - SAP HANA smart data integration and all its patches Product Availability Matrix (PAM) for SAP HANA SDI 2.0
> - Register Adapters with SAP Datasphere [page 119]

5.1 Preparing Data Provisioning Agent Connectivity

Most connection types that support creating views and accessing or replicating data via remote tables, for this purpose leverage SAP HANA Smart Data Integration (SDI) and its Data Provisioning Agent. Before using the connection, the agent requires an appropriate setup.

Context

The Data Provisioning Agent is a lightweight component running outside the SAP Datasphere environment. It hosts data provisioning adapters for connectivity to remote sources, enabling data federation and replication scenarios. The Data Provisioning Agent acts as a gateway to SAP Datasphere providing secure connectivity between the database of your SAP Datasphere tenant and the adapter-based remote sources. The Data

Provisioning Agent is managed by the Data Provisioning Server. It is required for all connections with SAP HANA smart data integration.

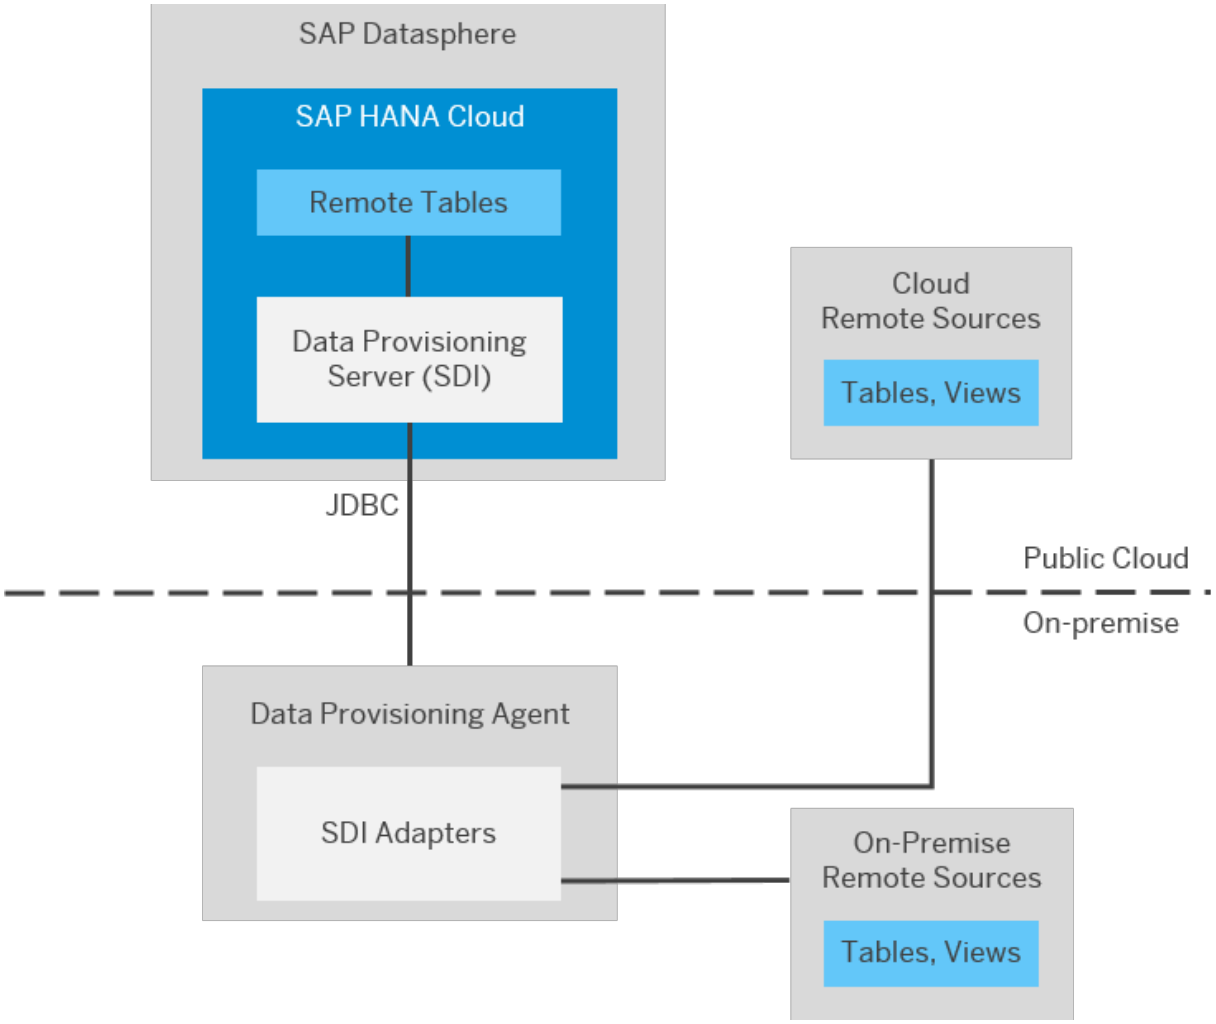

Through the Data Provisioning Agent, the preinstalled data provisioning adapters communicate with the Data Provisioning Server for connectivity, metadata browsing, and data access. The Data Provisioning Agent connects to SAP Datasphere using JDBC. It needs to be installed on a local host in your network and needs to be configured for use with SAP Datasphere.

For an overview of connection types that require a Data Provisioning Agent setup, see Preparing Connectivity for Connections [page 109].

> **Note**
>
> See also the guide Best Practices and Sizing Guide for Smart Data Integration (When used in SAP Datasphere) (published June 10, 2022) for information to consider when creating and using connections that are based on SDI and Data Provisioning Agent.

Procedure

To prepare connectivity via Data Provisioning Agent, perform the following steps:

1. Download and install the latest Data Provisioning Agent version on a host in your local network.

   > **Note**
   >
   > - We recommend to always use the latest released version of the Data Provisioning Agent. For information on supported and available versions for the Data Provisioning Agent, see the SAP HANA Smart Data Integration Product Availability Matrix (PAM).
   > - Make sure that all agents that you want to connect to SAP Datasphere have the same latest version.

   For more information, see Install the Data Provisioning Agent [page 115].

2. Add the external IPv4 address of the server on which your Data Provisioning Agent is running to the IP allowlist in SAP Datasphere. When using a proxy, the proxy's address needs to be included in IP allowlist as well.

   > **Note**
   >
   > For security reasons, all external connections to your SAP Datasphere instance are blocked by default. By adding external IPv4 addresses or address ranges to the allowlist you can manage external client connections.

   For more information, see Add IP address to IP Allowlist [page 128].

3. Connect the Data Provisioning Agent to SAP Datasphere.

   This includes configuring the agent and setting the user credentials in the agent.

   For more information, see Connect and Configure the Data Provisioning Agent [page 117].

4. Register the adapters with SAP Datasphere.

   > **Note**
   >
   > For third-party adapters, you need to download and install any necessary JDBC libraries before registering the adapters.

   For more information, see Register Adapters with SAP Datasphere [page 119].

Results

The registered adapters are available for creating connections to the supported remote sources and enabling these connections for creating views and accessing or replicating data via remote tables.

5.1.1 Install the Data Provisioning Agent

Download the latest Data Provisioning Agent 2.0 version from SAP Software Download Center and install it as a standalone installation on a Windows or Linux machine. If you have already installed an agent, check if

you need to update to the latest version. If you have more than one agent that you want to connect to SAP Datasphere, make sure to have the same latest version for all agents.

Context

Procedure

1. Plan and prepare the Data Provisioning Agent installation.
   a. Plan your installation to ensure that it meets your system landscape's needs.

      You can install the agent on any host system that has access to the sources you want to access, meets the minimum system requirements, and has any middleware required for source access installed. The agent should be installed on a host that you have full control over to view logs and restart, if necessary.

      For more information, see:

      - Planning and Preparation in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.
      - Supported Platforms and System Requirements in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.

   b. Download the latest Data Provisioning Agent *HANA DP AGENT 2.0* from the SAP Software Download Center.

      > **Note**
      >
      > - We recommend to always use the latest released version of the Data Provisioning Agent.
      > - Make sure that all agents that you want to connect to SAP Datasphere have the same latest version.
      > - Select your operating system before downloading the agent.

      For more information, see:

      - Software Download in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation
      - SAP HANA Smart Data Integration Product Availability Matrix (PAM) (for supported and available versions for the Data Provisioning Agent and operating system support)

2. Install the Data Provisioning Agent on a host in your local network.

   For more information, see Install from the Command Line in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.

   > **Note**
   >
   > If you have upgraded your Data Provisioning Agent to version 2.5.1 and want to create an Amazon Redshift connection, apply SAP note 2985825.

Related Information

Install the Data Provisioning Agent
Update the Data Provisioning Agent

5.1.2 Connect and Configure the Data Provisioning Agent

Connect the Data Provisioning Agent to the SAP HANA database of SAP Datasphere. This includes configuring the agent and setting the user credentials in the agent.

Procedure

1. In SAP Datasphere, register the Data Provisioning Agent.
   a. In the side navigation area, click (*System*) (*Configuration*) *Data Integration*.
   b. In the *On-Premise Agents* section, add a new tile to create a new agent registration in SAP Datasphere.
   c. In the following dialog, enter a unique name for your new agent registration.

      > **Note**
      >
      > The registration name cannot be changed later.

   d. Select *Create*.

      The *Agent Settings* dialog opens and provides you with information required to configure the Data Provisioning Agent on your local host:

      - Agent name
      - HANA server (host name)
      - HANA port
      - HANA user name for agent messaging
      - HANA user password for agent messaging

      > **Note**
      >
      > Either keep the *Agent Settings* dialog open, or note down the information before closing the dialog.

2. At the command line, connect the agent to SAP HANA using JDBC. Perform the following steps:
   a. Navigate to `<DPAgent_root>/bin/`. `<DPAgent_root>` is the Data Provisioning Agent installation root location. By default, on Windows, this is `C:\usr\sap\dataprovagent`, and on Linux it is `/usr/sap/dataprovagent`.
   b. Start the agent using the following command:

      On Linux: `./ dpagent_servicedaemon.sh start`

      On Windows: `dpagent_servicedaemon_start.bat`
   c. Start the command-line agent configuration tool using the following command:

On Linux:

```
<DPAgent_root>/bin/agentcli.sh --configAgent
```

On Windows:

```
<DPAgent_root>/bin/agentcli.bat --configAgent
```

d. Choose *SAP HANA Connection*.
e. Choose *Connect to SAP Data Warehouse Cloud via JDBC*.
f. Enter the name of the agent registration (agent name).
g. Enter **`true`** to use an encrypted connection over JDBC.

> → Tip
>
> An encrypted connection is always required when connecting to SAP HANA in a cloud-based environment.

h. Enter the host name (HANA server) and port number (HANA port) for the SAP Datasphere instance.

   For example:

   - Host name: **`<instance_name>.hanacloud.ondemand.com`**
   - Port number: **`443`**

i. If HTTPS traffic from your agent host is routed through a proxy, enter **`true`** and specify any required proxy information as prompted.
   1. Enter **`true`** to specify that the proxy is an HTTP proxy.
   2. Enter the proxy host and port.
   3. If you use proxy authentication, enter **`true`** and provide a proxy user name and password.
j. Enter the credentials for the HANA user for agent messaging.

   The HANA user for agent messaging is used only for messaging between the agent and SAP Datasphere.
k. Confirm that you want to save the connection settings you have made by entering **`true`**.

> ⓘ Note
>
> Any existing agent connection settings will be overwritten.

l. Stop and restart the Data Provisioning Agent.

   On Linux:

```
<DPAgent_root>/bin/agentcli.sh --configAgent
```

   On Windows:

```
<DPAgent_root>/bin/agentcli.bat --configAgent
```

   1. To stop the agent, choose *Start or Stop Agent*, and then choose *Stop Agent*.
   2. Choose *Start Agent* to restart the agent.
   3. Choose *Agent Status* to check the connection status. If the connection succeeded, you should see `Agent connected to HANA: Yes`.
   4. Choose *Quit* to exit the script.

3. In SAP Datasphere, if you have kept the *Agent Settings* dialog open, you can now close it.

Results

The Data Provisioning Agent is now connected.

If the tile of the registered Data Provisioning Agent doesn't display the updated connection status, select *Refresh Agents*.

Related Information

Troubleshooting the Data Provisioning Agent (SAP HANA Smart Data Integration) [page 168]

5.1.3 Register Adapters with SAP Datasphere

After configuring the Data Provisioning Agent, in SAP Datasphere, register the Data Provisioning adapters that are needed to connect to on-premise sources.

Prerequisites

For third-party adapters, ensure that you have downloaded and installed any necessary JDBC libraries. Place the files in the `<DPAgent_root>/lib` folder before registering the adapters with SAP Datasphere. For connection types *Amazon Redshift* and *Generic JDBC*, place the file in the `<DPAgent_root>/camel/lib` folder. For information about the proper JDBC library for your source, see the *SAP HANA smart data integration Product Availability Matrix (PAM)*. Search for the library in the internet and download it from an appropriate web page.

Procedure

1. In the side navigation area, click (*System*) (*Configuration*) *Data Integration*.
2. In the *On-Premise Agents* section, click the *Adapters* button to display the agents with their adapter information.
3. Click *(menu)* and then *Edit*.
4. In the *Agent Settings* dialog, under *Agent Adapters* select the adapters.
5. Close the dialog.
   The selected adapters are registered with SAP Datasphere and available for creating connections to the supported on-premise sources.

Next Steps

To use new functionality of an already registered adapter or to update the adapter in case of issues that have been fixed in a new agent version, you can refresh the adapter by clicking the ••• *(menu)* button and then choosing *Refresh*.

Related Information

SAP HANA smart data integration and all its patches Product Availability Matrix (PAM) for SAP HANA SDI 2.0

5.1.4 Prerequisites for ABAP RFC Streaming

If you want to stream ABAP tables for loading large amounts of data without running into memory issues it is required to meet the following requirements.

- You need to create an RFC destination in the ABAP source system. With the RFC destination you register the Data Provisioning agent as server program in the source system.
  Using transaction SM59, you create a TCP/IP connection with a user-defined name. The connection should be created with "Registered Server Program" as "Activation Type". Specify "IM_HANA_ABAPADAPTER_*" as a filter for the "Program ID" field, or leave it empty.
- Successful registration on an SAP Gateway requires that suitable security privileges are configured. For example:
  - Set up an Access Control List (ACL) that controls which host can connect to the gateway. That file should contain something similar to the following syntax: `<permit> <ip-address[/mask]> [tracelevel] [# comment]`. `<ip-address>` here is the IP of the server on which Data Provisioning agent has been installed.
    For more information, see the *Gateway* documentation in the SAP help for your source system version, for example Configuring Network-Based Access Control Lists (ACL) in the *SAP NetWeaver 7.5* documentation.
  - You may also want to configure a `reginfo` file to control permissions to register external programs.

5.2 Preparing Cloud Connector Connectivity

Connections to on-premise sources used for data flows, replication flows, and other use cases require Cloud Connector to act as link between SAP Datasphere and the source. Before creating the connection, the Cloud Connector requires an appropriate setup.

Context

Cloud Connector serves as a link between SAP Datasphere and your on-premise sources and is required for connections that you want to use for:

- Data flows
- Replication flows
- Model import from:
  - SAP BW/4HANA Model Transfer connections (Cloud Connector is required for the live data connection of type tunnel that you need to create the model import connection)
  - SAP S/4HANA On-Premise connections (Cloud Connector is required for the live data connection of type tunnel that you need to search for the entities in the SAP S/4HANA system)
- Remote tables (only for SAP HANA on-premise via SAP HANA Smart Data Access)

For an overview of connection types that require a Cloud Connector setup to be able to use any of these features, see Preparing Connectivity for Connections [page 109].

Procedure

To prepare connectivity via Cloud Connector, perform the following steps:

1. Install the Cloud Connector in your on-premise network.
   For more information, see Cloud Connector Installation in the *SAP BTP Connectivity* documentation.
2. Make sure to hold the SAP Datasphere subaccount information ready. You can find the information in ⓘ (*System*) ▸ ⚙ (*Administration*) ▸ *Data Source Configuration*.
   For more information, see Set Up Cloud Connector in SAP Datasphere [page 127].
3. In the Cloud Connector administration, set up and configure Cloud Connector according to your requirements.
   For more information, see Configure Cloud Connector [page 122].
4. If you have connected multiple Cloud Connector instances to your subaccount and you want to use these locations for your connections, add the location IDs in ⓘ (*System*) ▸ ⚙ (*Administration*) ▸ *Data Source Configuration*.
   For more information, see Set Up Cloud Connector in SAP Datasphere [page 127].
5. If you you want to create SAP BW/4HANA Model Transfer connections or SAP S/4HANA On-Premise connections for model import, make sure you have switched on *Allow live data to securely leave my network* in ⓘ (*System*) ▸ ⚙ (*Administration*) ▸ *Data Source Configuration*.
   For more information, see Set Up Cloud Connector in SAP Datasphere [page 127].

Result

The Cloud Connector respectively Cloud Connector instances are available for creating connections and enabling these for the supported features.

Related Links

Frequently Asked Questions (about the Cloud Connector) in the *SAP BTP Connectivity* documentation

5.2.1 Configure Cloud Connector

Configure Cloud Connector before connecting to on-premise sources and using them in various use cases. In the Cloud Connector administation, connect the SAP Datasphere subaccount to your Cloud Connector, add a mapping to each relevant source system in your network, and specify accessible resources for each source system.

Prerequisites

Before configuring the Cloud Connector, the following prerequisites must be fulfilled:

- The Cloud Connector is installed in your on-premise network.
  For more information, see Cloud Connector Installation in the *SAP BTP Connectivity* documentation.
- If you are using egress firewalling, add the following domains (wildcard) to the firewall/proxy allowlist in your on-premise network:
  - `*.hanacloud.ondemand.com`
  - `*.k8s-hana.ondemand.com`
- Before configuring the Cloud Connector, you or the owner of your organisation will need an SAP Business Technology Platform (SAP BTP) account. If you don't have an account yet, create an account by clicking *Register* in the SAP BTP cockpit.
- During Cloud Connector configuration you will need information for your SAP Datasphere subaccount.
  Make sure that you have the subaccount information available in *System* > *Administration* > *Data Source Configuration* > *SAP BTP Core Account*.
  For more information, see Set Up Cloud Connector in SAP Datasphere [page 127].

Context

For more information about the supported use cases depending on the connection type, see Preparing Cloud Connector Connectivity [page 121].

Procedure

1. Log on to the *Cloud Connector Administration* on `https://<hostname>:8443`.

   `<hostname>` refers to the machine on which the Cloud Connector is installed. If installed on your machine, you can simply enter localhost.

2. To connect the SAP Datasphere subaccount to your Cloud Connector, perform the following steps:
   a. In the side navigation area of the *Cloud Connector Administration*, click *Connector* to open the *Connector* page and click + *Add Subaccount* to open the *Add Subaccount* dialog.
   b. Enter or select the following information to add the SAP Datasphere subaccount to the Cloud Connector.

> **Note**
>
> You can find the subaccount, region, and subaccount user information in SAP Datasphere under *System* > *Administration* > *Data Source Configuration* > *SAP BTP Core Account* > *Account Information*.

Property	*Description*
Region	Select your region host from the list.
Subaccount	Add your SAP Datasphere subaccount name.
Display Name	[optional] Add a name for the account.
Subaccount User	Add your subaccount (S-User) username.
Password	Add your S-User password for the SAP Business Technology Platform.
Location ID	[optional] Define a location ID that identifies the location of this Cloud Connector for the subaccount. **Note** • Using location IDs you can connect multiple Cloud Connector instances to your subaccount. If you don't specify any value, the default is used. For more information, see Managing Subaccounts in the *SAP BTP Connectivity* documentation. • Each Cloud Connector instance must use a different location, and an error will appear if you choose a location that is already been used. • We recommend that you leave the *Location ID* empty if you don't plan to set up multiple Cloud Connectors in your system landscape.
Description	(Optional) Add a description for the Cloud Connector.

   c. Click *Save*.

In the *Subaccount Dashboard* section of the *Connector* page, you can see all subaccounts added to the Cloud Connector at a glance. After you added your subaccount, you can check the status to verify that the Cloud Connector is connected to the subaccount.

3. To allow SAP Datasphere to access systems (on-premise) in your network, you must specify the systems and the accessible resources in the Cloud Connector (URL paths or function module names depending on the used protocol). Perform the following steps for each system that you want to be made available by the Cloud Connector:

   a. In the side navigation area, under your subaccount menu, click *Cloud To On-Premise* and then + *(Add)*in the *Mapping Virtual To Internal System* section of the *Access Control* tab to open the *Add System Mapping* dialog.

   > ### Note
   >
   > The side navigation area shows the display name of your subaccount. If the area shows another subaccount, select your subaccount from the *Subaccount* field of the *Cloud Connector Administration*.

   b. Add your system mapping information to configure access control and save your configuration.

   The procedure to add your system mapping information is specific to the protocol that you are using for communication. The relevant protocols are:

Connection Type (Feature used with the Connection)	Protocol
SAP ABAP (data flows, replication flows) *SAP BW* (data flows) *SAP ECC* (data flows) *SAP S/4HANA On-Premise* (data flows, replication flows, model import)	RFC
SAP BW/4HANA Model Transfer (model import) *SAP S/4HANA On-Premise* (model import)	HTTPS
SAP HANA on-premise only (data flows, replication flows, remote tables via SAP HANA Smart Data Access and Cloud Connector)	TCP For information about how to enable encrypted communication, see the *Security* properties in Configuring Connection Properties (SAP HANA on-premise).
Generic OData (data flows)	HTTPS
Microsoft SQL Server (data flows)	TCP
Apache Kafka on-premise only (replication flows)	TCP
Confluent - Confluent Platform on-premise only (replication flows)	• For the Kafka broker: TCP • For the Schema Registry: HTTP

   For more information, see Configure Access Control (HTTP) and Configure Access Control (RFC) in the *SAP BTP Connectivity* documentation.

> **Note**
>
> When adding the system mapping information, you enter internal and virtual system information. The internal host and port specify the actual host and port under which the backend system can be reached within the intranet. It must be an existing network address that can be resolved on the intranet and has network visibility for the Cloud Connector. The Cloud Connector tries to forward the request to the network address specified by the internal host and port, so this address needs to be real. The virtual host name and port represent the fully qualified domain name of the related system in the cloud.
>
> We recommend to use a virtual (cloud-side) name that is different from the internal name.

c. To grant access only to the resources needed by SAP Datasphere, select the system host you just added from the *Mapping Virtual To Internal System* list, and for each resource that you want to allow to be invoked on that host click + *(Add)* in the *Resources Of* section to open the *Add Resource* dialog.

d. Depending on the connection type, protocol, and use case, add the required resources:

Connection Type	Resource Type (depending on protocol)	Resources
SAP BW/4HANA Model Import	*URL Path* (for HTTPS)	• `/sap/opu/odata/sap/ESH_SEARCH_SRV/SearchQueries` • `/sap/bw4/v1/dwc/dbinfo` • `/sap/bw4/v1/dwc/metadata/queryviews` – *Path and all sub-paths* • `/sap/bw4/v1/dwc/metadata/treestructure` – *Path and all sub-paths* • `/sap/bw/ina` – *Path and all sub-paths*
SAP S/4HANA On-Premise	*URL Path* (for HTTPS)	`/`

Connection Type	Resource Type (depending on protocol)	Resources
SAP ABAP *SAP S/4HANA On-Premise*	*Function Name* (name of the function module for RFC)	For accessing data using CDS view extraction: • DHAMB_ – *Prefix* • DHAPE_ – *Prefix* • RFC_FUNCTION_SEARCH For accessing data based on tables with SAP LT Replication Server: • LTAMB_ – *Prefix* • LTAPE_ – *Prefix* • RFC_FUNCTION_SEARCH
SAP BW *SAP ECC*	*Function Name* (name of the function module for RFC)	For accessing data using ODP connectivity (for legacy systems that do not have the ABAP Pipeline Engine extension or DMIS Addon installed): • /SAPDS/ – *Prefix* • RFC_FUNCTION_SEARCH • RODPS_REPL_ – *Prefix*
Confluent - Confluent Platform (for the Schema Registry)	*URL Path* (for HTTPS)	/

For more information, see Configure Access Control (HTTP) and Configure Access Control (RFC) in the *SAP BTP Connectivity* documentation.

e. Choose *Save*.

4. [optional] To enable secure network communication (SNC) for data flows, configure SNC in the Cloud Connector.

   For more information, see Initial Configuration (RFC) in the *SAP BTP Connectivity* documentation.

Next Steps

1. If you have defined a location ID in the Cloud Connector configuration and want to use it when creating connections, you need to add the location ID in ⓘ (*System*) > ⚙ (*Administration*) > *Data Source Configuration*.
   For more information, see Set Up Cloud Connector in SAP Datasphere [page 127].
2. If you want to create SAP BW/4HANA Model Transfer connections or SAP S/4HANA On-Premise connections for model import, you need to switch on *Allow live data to securely leave my network* in ⓘ (*System*) > ⚙ (*Administration*) > *Data Source Configuration*
   For more information, see Set Up Cloud Connector in SAP Datasphere [page 127].

You can now create your connections in SAP Datasphere.

Related Information

For answers to the most common questions about the Cloud Connector, see Frequently Asked Questions in the *SAP BTP Connectivity* documentation.

5.2.2 Set Up Cloud Connector in SAP Datasphere

Receive SAP Datasphere subaccount information required for Cloud Connector configuration and complete Cloud Connector setup for creating SAP BW/4HANA Model Transfer connections and for using multiple Cloud Connector instances.

Context

The Cloud Connector allows you to connect to on-premise data sources and use them in various use cases depending on the connection type.

For more information, see Preparing Cloud Connector Connectivity [page 121].

Procedure

1. In the side navigation area, click (*System*) > (*Administration*) > *Data Source Configuration*.

   > **Note**
   >
   > If your tenant was provisioned prior to version 2021.03, click *(Product Switch)* > *Analytics* > *System* > *Administration* > *Data Source Configuration*.

2. Perform the required tasks:
   - Receive the SAP Datasphere subaccount information that is required during Cloud Connector configuration.
     To receive the SAP Datasphere subaccount information, the subaccount needs to be linked to the user ID of your SAP BTP account. In the *SAP BTP Core Account* section, you can check if this has been done and the information is already available in *Account Information*.
     During Cloud Connector configuration, you will then need to enter the following information from your SAP Datasphere subaccount:
     - *Subaccount*
     - *Region Host*
     - *Subaccount User*

     If you have an account but cannot see the *Account Information*, enter the SAP BTP user ID. This ID is typically the email address you used to create your SAP BTP account. After you have entered the ID you can see the *Account Information* for SAP Datasphere:

> **Note**
>
> If you don't have an SAP Business Technology Platform (SAP BTP) user account yet, create an account in the SAP BTP cockpit by clicking *Register* in the cockpit.

- To be able to use the Cloud Connector for SAP BW/4HANA Model Transfer connections to import analytic queries with the *Model Transfer Wizard* and for SAP S/4HANA On-Premise connections to import ABAP CDS Views with the *Import Entities* wizard, switch on *Allow live data to securely leave my network* in the *Live Data Sources* section.

  > **Note**
  >
  > The *Allow live data to securely leave my network* switch is audited, so that administrators can see who switched this feature on and off. To see the changes in the switch state, go to (*Security*) (*Activities*), and search for **`ALLOW_LIVE_DATA_MOVEMENT`**.

- If you have connected multiple Cloud Connector instances to your subaccount with different location IDs and you want to offer them for selection when creating connections using a Cloud Connector, in the *On-premise data sources* section, add the appropriate location IDs. If you don't add any location IDs here, the default location will be used.
  Cloud Connector location IDs identify Cloud Connector instances that are deployed in various locations of a customer's premises and connected to the same subaccount. Starting with Cloud Connector 2.9.0, it is possible to connect multiple Cloud Connectors to a subaccount as long as their location ID is different.

5.3 Add IP address to IP Allowlist

Clients in your local network need an entry in the appropriate IP allowlist in SAP Datasphere. Cloud Connectors in your local network only require an entry if you want to use them for federation and replication with remote tables from on-premise systems.

Context

To secure your environment, you can control the range of IPv4 addresses that get access to the database of your SAP Datasphere by adding them to an allowlist.

You need to provide the **external (public) IPv4 address (range)** of the client directly connecting to the database of SAP Datasphere. This client might be an SAP HANA smart data integration Data Provisioning Agent on a server, a 3rd party ETL or analytics tool, or any other JDBC-client. If you're using a network firewall with a proxy, you need to provide the public IPv4 address of your proxy.

Internet Protocol version 4 addresses (IPv4 addresses) have a size of 32 bits and are represented in dot-decimal notation, *192.168.100.1* for example. The external IPv4 address is the address that the internet and computers outside your local network can use to identify your system.

The address can either be a single IPv4 address or a range specified with a Classless Inter-Domain Routing suffix (CIDR suffix). An example for a CIDR suffix is **/24** which represents 256 addresses and is typically used

for a large local area network (LAN). The CIDR notation for the IPv4 address above would be: *192.168.100.1/24* to denote the IP addresses between *192.168.100.0* and *192.168.100.255* (the leftmost 24 bits of the address in binary notation are fixed). The external (public) IP address (range) to enter into the allowlist will be outside of the range *192.168.0.0/16*. You can find more information on Classless Inter-Domain Routing on Wikipedia .

> **Note**
>
> The number of entries in the allowlist is limited. Once the limit has been reached, you won't be able to add entries. Therefore, please consider which IP addresses should be added and whether the number of allowlist entries can be reduced by using ranges to request as few allowlist entries as possible.

Procedure

1. In the side navigation area, click (*System*) (*Configuration*) *IP Allowlist* .
2. From the *IP Allowlist* dropdown, select the appropriate list:
   - *Trusted IPs*: For clients such as an Data Provisioning Agent on a server, 3rd party ETL or analytics tools, or any other JDBC-client
   - *Trusted Cloud Connector IPs*: For Cloud Connectors that you want to use for federation and replication with remote tables from on-premise systems such as SAP HANA

   The selected list shows all IP addresses that are allowed to connect to the SAP Datasphere database.
3. Click *Add* to open the *Allow IP Addresses* dialog.

   > **Note**
   >
   > Once the number of entries in the allowlist has reached its limit, the *Add* button will be disabled.

4. In the *CIDR* field of the dialog, either provide a single IPv4 address or a range specified with a CIDR suffix.

   > **Note**
   >
   > Please make sure that you provide the **external** IPv4 address of your client respectively proxy when using a network firewall. The IP you enter needs to be your public internet IP.

5. [optional] You can add a description of up to 120 characters to better understand your IP entries.
6. In the dialog, click *Add* to return to the list.
7. To save your newly added IP to the allowlist on the database, click *Save* in the pushbutton bar of your list.

   > **Note**
   >
   > Updating the allowlist in the database requires some time. To check if your changes have been applied, click *Refresh*.

Next Steps

You can also select and edit an entry from the list if an IP address has changed, or you can delete IPs if they are not required anymore to prevent them from accessing the database of SAP Datasphere. To update the allowlist

in the database with any change you made, click *Save* and be reminded that the update in the database might take some time.

5.4 Finding SAP Datasphere IP addresses

Find externally facing IP addresses that for particular remote applications must be added to allowlists before you can to use connections to these remote applications.

Particular remote applications or sources that you might want to access with SAP Datasphere restrict access to their instances and require external SAP Datasphere IP address information to be added to an allowlist in the remote application before first trying to access the application .

Outbound IP Address

The network for Amazon Redshift, Microsoft Azure SQL Database, or SAP SuccessFactors instances, for example, is protected by a firewall that controls incoming traffic. To be able to use connections with these connection types for data flows (and Microsoft Azure SQL Database also for replication flows), the connected sources require the SAP Datasphere outbound IP address to be added to an allowlist.

Find the *Outbound IP Address* in the last step of the connection creation wizard.

Administrators can find the *Outbound IP Address* from the side navigation area by clicking ⓘ (*System*) ⓘ (*About*) and expanding the *More* section in the dialog.

HANA IP Addresses

Access to SAP SuccessFactors instances is restricted. To be able to use a *SAP SuccessFactors* connection for remote tables and view building, the connected source requires the externally facing IP addresses of theSAP Datasphere tenant to be added to an allowlist.

Administrators can find the *HANA IP Addresses* from the side navigation area by clicking ⓘ (*System*) ⓘ (*About*) and expanding the *More* section in the dialog.

For more information about adding the IP addresses in SAP SuccessFactors, see Adding an IP Restriction in the *SAP SuccessFactors platform* documentation.

5.5 Manage Certificates for Connections

To import a certificate into the SAP Datasphere trust chain, obtain the certificate from the target endpoint and upload it to SAP Datasphere.

Prerequisites

You have downloaded the required SSL/TLS certificate from an appropriate website. As one option for downloading, common browsers provide functionality to export these certificates.

> **Note**
>
> - Only X.509 Base64-encoded certificates enclosed between "-----BEGIN CERTIFICATE-----" and "-----END CERTIFICATE-----" are supported. The common filename extension for the certificates is .pem (Privacy-enhanced Electronic Mail). We also support filename extensions .crt and .cer.
> - A certificate used in one region might differ from those used in other regions. Also, some sources, such as Amazon Athena, might require more than one certificate.
> - Remember that all certificates can expire.
> - If you have a problem with a certificate, please contact your cloud company for assistance.

Context

For connections secured by leveraging HTTPS as the underlying transport protocol (using SSL/TLS transport encryption), the server certificate must be trusted.

> **Note**
>
> You can create connections to remote systems which require a certificate upload without having uploaded the necessary certificate. Validating a connection without valid server certificate will fail though, and you won't be able to use the connection.

Procedure

1. In the side navigation area, click (*System*) > (*Configuration*) > *Security*.
2. Click + *Add Certificate*.
3. In the *Upload Certificate* dialog, browse your local directory and select the certificate.
4. Enter a description to provide intelligible information on the certificate, for example to point out to which connection type the certificate applies.
5. Choose *Upload*.

Results

In the overview, you can see the certificate with its creation and expiry date. From the overview, you can delete certificates if required.

5.6 Upload Third-Party ODBC Drivers (Required for Data Flows)

To enable access to a non-SAP database via ODBC to use it as a source for data flows, you need to upload the required ODBC driver files to SAP Datasphere.

Prerequisites

- Search for the required driver files in the internet, make sure you have selected the correct driver files (identified by their SHA256-formatted fingerprint) and download them from an appropriate web page (see below).
- Ensure you have a valid license for the driver files.

Context

Drivers are required for the following connection types (if several driver versions are supported, we recommend to use the newest supported version mentioned below):

Connection Type	Driver to be uploaded	SHA256 Fingerprint	Download Site
Amazon Redshift Connections	AmazonRedshiftODBC-64-bit-1.4.11.1000-1.x86_64.rpm	6d811e2f198a030274bf9f099d4c828b1b071b78e99432eee1531d4988768a22	`https://docs.aws.amazon.com`
	AmazonRedshiftODBC-64-bit-1.4.65.1000-1.x86_64.rpm	ee79a8d41760a90b6fa2e1a074e33b0518e3393afd305f0bee843b5393e10df0	

Connection Type	Driver to be uploaded	SHA256 Fingerprint	Download Site
Oracle Connections	instantclient-basiclite-linux.x64-12.2.0.1.0.zip **Note** Make sure to select the *Basic Light* package zip file. The package applies to version 12c supported by the Oracle connection type.	1c3adb36f9605aae84ae98461bd6ee9eb26b303cace3f5534cd7985d470d0494	`https://oracle.com`

Connection Type	Driver to be uploaded	SHA256 Fingerprint	Download Site

	instantclient-basiclite-linux.x64-19.17.0.0.0dbru.zip
	
	Note
	• Make sure to select the *Basic Light* package zip file. The package applies to all versions supported by the Oracle connection type (Oracle 12c, Oracle 18c, and Oracle 19c).
	• Additional files are required if SSL is used:
	  • oraclepki.jar (SHA256 fingerprint: e408e7ae67650917dbce3ad263829bdc6c791d50d4db2fd59aeeb5503175499b)
	  • package zip file. The package applies to allosdt_cert.jar (SHA256 fingerprint: 6b152d4332bd39f258a88e58b9215a926048d740e148971fe1628b09060176a8)
	  • osdt_core.jar (SHA256 fingerprint: c25e30184bb94c6da1227c8256f0e1336acb97b29229edb4a
		ea4a9557c6355f5b56b648b7dff47db79a1403b7e9f7abeca9e1a0e952498e13	Driver: `https://download.oracle.com/otn_software/linux/instantclient/1917000/instantclient-basiclite-linux.x64-19.17.0.0.0dbru.zip`
			
			Additional files if SSL is used:
			• `https://repo1.maven.org/maven2/com/oracle/database/security/oraclepki/19.17.0.0/oraclepki-19.17.0.0.jar`
			• `https://repo1.maven.org/maven2/com/oracle/database/security/osdt_core/19.17.0.0/osdt_core-19.17.0.0.jar`
			• `https://repo1.maven.org/maven2/com/oracle/database/security/osdt_cert/19.17.0.0/osdt_cert-19.17.0.0.jar`

Connection Type	Driver to be uploaded	SHA256 Fingerprint	Download Site
	acf27167b96075e)		
Google BigQuery Connections	SimbaODBCDriverforGoogleBigQuery_2.3.1.1001-Linux.tar.gz	abf4551d621c26f4fa30539e7ece2a47daaf6e1d67c59e5b7e79c43a3335018f	`https://storage.googleapis.com/simba-bq-release/odbc/SimbaODBCDriverforGoogleBigQuery_2.3.1.1001-Linux.tar.gz`
	SimbaODBCDriverforGoogleBigQuery_3.0.0.1001-Linux.tar.gz	58d3c9acfb93f0d26c081a230ff664a16c8544d567792ebc5436beb31e9e28e4	`https://cloud.google.com/bigquery/providers/simba-drivers`

When uploading the drivers, they are identified by their SHA256-formatted fingerprint. You can verify the fingerprint with the following command:

- Windows 10: In PowerShell, run the following command:
  `Get-Filehash <driver file> -Algorithm SHA256`
- Linux/MacOS: In a unix-compliant shell, run the following command:
  `shasum -a 256 <driver file>`

Upload a Driver

Perform the following steps before creating the first Amazon Redshift, Oracle, or Google BigQuery connection that you want to use for data flows.

1. In the side navigation area, click (*System*) (*Configuration*) *Data Integration*.
2. Go to *Third-Party Drivers* and choose + *Upload*.
3. In the following dialog box, choose *Browse* to select the driver file from your download location.

   > **Note**
   >
   > The fingerprint of the driver file name to be uploaded must match the fingerprint mentioned above.

4. Choose *Upload*.
5. Choose *sync* to synchronize the driver with the underlying component. Wait for about 5 to 10 minutes to finish synchronization before you start creating connections or using data flows with the connection.

Remove (and Re-Upload) a Driver

You might need to remove a driver when you want to upload a new version of the driver or your licence agreement has terminated.

1. Select the driver and choose *Delete*.
2. If you're using a connection that requires the removed driver for data flows, choose *Upload* to re-upload the driver to make sure that you can continue using the data flows.
3. Choose *sync* to synchronize the driver changes with the underlying component. Once the synchronization has finished, you can continue using data flows with the connection, or, if you haven't uploaded a new driver, you won't be able to use data flows with the connection anymore unless you re-upload the driver.

Troubleshooting

If a data flow fails with the error message saying that the driver could not be found, check that the drivers are uploaded and start synchronization.

5.7 Prepare Connectivity to Adverity

To be able to successfully validate and use a connection to Adverity for view building certain preparations have to be made.

Before you can use the connection, the following is required:

- In an Adverity workspace, you have prepared a datastream that connects to the data source for which you want to create the connection.
- In SAP Datasphere, you have added the necessary Adverity IP addresses to the IP allowlist. For more information, see Add IP address to IP Allowlist [page 128].

> **Note**
>
> To get the relevant IP addresses, please contact your Adverity Account Manager or the Adverity Support team.

5.8 Prepare Connectivity to Amazon Athena

To be able to successfully validate and use a connection to Amazon Athena for remote tables certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- A DW administrator has uploaded the server certificates to SAP Datasphere. Two certificates are required, one for Amazon Athena and one for Amazon S3. Region-specific certificates might be required for Amazon Athena.
  For more information, see Manage Certificates for Connections [page 131].

5.9 Prepare Connectivity to Apache Kafka

To be able to use a connection to Apache Kafka (on-premise) for replication flows, certain preparations have to be made.

Replication Flows

Before you can use the connection for replication flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to the Apache Kafka on-premise implementation.

5.10 Prepare Connectivity to Confluent

To be able to use a connection to Confluent Platform (on-premise) for replication flows, certain preparations have to be made.

Replication Flows

Before you can use the connection for replication flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to Confluent Platform (Kafka brokers) and to the Schema Registry.

  > ⓘ Note
  >
  > Separate Cloud Connector instances might be used for the two endpoints. The Schema Registry might be used in one Cloud Connector location is while connecting to the Kafka brokers happens in another location.

5.11 Prepare Connectivity to Amazon Redshift

To be able to successfully validate and use a connection to an Amazon Redshift database for remote tables or data flows certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the CamelJdbcAdapter.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- An administrator has downloadad and installed the required JDBC library in the `<DPAgent_root>/camel/lib` folder and restarted the Data Provisioning Agent before registering the adapter with SAP Datasphere.

Data Flows

Before you can use the connection for data flows, the following is required:

- The outbound IP has been added to the source allowlist.
  For information on where a DW administrator can find the IP address, see Finding SAP Datasphere IP addresses [page 130].
- A DW administrator has uploaded the required ODBC driver file to SAP Datasphere.
  For more information, see Upload Third-Party ODBC Drivers (Required for Data Flows) [page 132].

5.12 Prepare Connectivity for Cloud Data Integration

To be able to successfully validate and use a Cloud Data Integration connection for remote tables or data flows certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the CloudDataIntegrationAdapter.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- For ABAP-based cloud SAP systems such as SAP S/4HANA Cloud or SAP Marketing Cloud: A communication arrangement has been created for communication scenario `SAP_COM_0531` in the source system.
  For more information, see Integrating CDI in the *SAP S/4HANA Cloud* documentation.

Data Flows

Before you can use the connection for data flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to your on-premise source.
  For more information, see Configure Cloud Connector [page 122].
- For ABAP-based cloud SAP systems such as SAP S/4HANA Cloud or SAP Marketing Cloud: A communication arrangement has been created for communication scenario `SAP_COM_0531` in the source system.
  For more information, see Integrating CDI in the *SAP S/4HANA Cloud* documentation.

5.13 Prepare Connectivity for Generic JDBC

To be able to successfully validate and use a Generic JDBC connection for remote tables certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the CamelJdbcAdapter.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- It has been checked that the data source is supported by the CamelJdbcAdapter.
  For latest information about supported data sources and versions, see the SAP HANA Smart Data Integration Product Availability Matrix (PAM).

  > **Note**
  >
  > For information about unsupported data sources, see 3130999.

- An administrator has downloadad and installed the required JDBC library in the `<DPAgent_root>/camel/lib` folder and restarted the Data Provisioning Agent before registering the adapter with SAP Datasphere.
  For more information, see Set up the Camel JDBC Adapter in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality Installation and Configuration Guide*.
  For information about the proper JDBC library for your source, see the *SAP HANA smart data integration Product Availability Matrix (PAM)*.

5.14 Prepare Connectivity for Generic OData

To be able to successfully validate and use a connection to an OData service for remote tables or data flows certain preparations have to be made.

General

Before you can use the connection, the following is required:

- The OData service URL needs to be publicly available.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- A DW administrator has uploaded the server certificate to SAP Datasphere.
  For more information, see Manage Certificates for Connections [page 131].

Data Flows

Before you can use the connection for data flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to your on-premise source.
  For more information, see Configure Cloud Connector [page 122].

5.15 Prepare Connectivity for Generic SFTP

To create a Generic SFTP connection, the host's public key is required.

The host's public key should be provided through a trustable channel. If your Windows 10, Linux, or MacOS machine has such a channel, follow the steps below by replacing any occurence of `$HOST` with the host value of your connection, and `$PORT` with the port value.

Use the resulting file `host_key.pub.txt` (found in the directory where you run the following command) to upload the *Host Key* when creating your connection.

- Windows 10: In PowerShell, run the following command:
  ```
  (ssh-keyscan -t rsa -p $PORT $HOST 2>$null) -replace '^[^ ]* ','' >
  host_key.pub.txt
  ```
- Linux/MacOS: In a unix-compliant shell with both `ssh-keyscan` and `sed` commands (both are usually already installed in your system), obtain the key through the following command:
  ```
  ssh-keyscan -t rsa -p $PORT $HOST 2>/dev/null | sed "s/^[^ ]* //" >
  host_key.pub.txt
  ```

> **Note**
>
> If your machine doesn't have a trustable channel, we recommend asking your administrator for the public host key to avoid man-in-the-middle attacks.

5.16 Prepare Connectivity to Google BigQuery

To be able to successfully validate and use a connection to a Google BigQuery data source for remote tables, certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- A DW administrator has uploaded the server certificate to SAP Datasphere.
  For more information, see Manage Certificates for Connections [page 131].

Data Flows and Replication Flows

Before you can use the connection for data flows, the following is required:

- A DW administrator has uploaded the required ODBC driver file to SAP Datasphere.
  For more information, see Upload Third-Party ODBC Drivers (Required for Data Flows) [page 132].

5.17 Prepare Connectivity to Microsoft Azure SQL Database

To be able to successfully validate and use a connection to Microsoft Azure SQL database for remote tables or data flows certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the MssqlLogReaderAdapter.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- An administrator has downloadad and installed the required JDBC library in the `<DPAgent_root>/lib` folder before registering the adapter with SAP Datasphere.
- To use Microsoft SQL Server trigger-based replication, the user entered in the connection credentials needs to have the required privileges and permissions. For more information, see Required Permissions for SQL Server Trigger-Based Replication in the *Installation and Configuration Guide* for *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality*

Data Flows and Replication Flows

Before you can use the connection for data flows and replication flows, the following is required:

- The outbound IP has been added to the source allowlist.
  For information on where a DW administrator can find the IP address, see Finding SAP Datasphere IP addresses [page 130].

5.18 Prepare Connectivity to Microsoft SQL Server

To be able to successfully validate and use a connection to a Microsoft SQL Server for remote tables or data flows, certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the MssqlLogReaderAdapter.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- An administrator has downloadad and installed the required JDBC library in the `<DPAgent_root>/lib` folder before registering the adapter with SAP Datasphere.
- Required Permissions for SQL Server Trigger-Based Replication in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality Installation and Configuration Guide*

Data Flows

Before you can use the connection for data flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to your on-premise source.
  For more information, see Configure Cloud Connector [page 122].

  > **Note**
  >
  > Cloud Connector is not required if your Microsoft SQL Server database is available on the public internet.

- The required driver is pre-bundled and doesn't need to be uploaded by an administrator.

5.19 Prepare Connectivity to SAP Open Connectors

Integrate SAP Open Connectors with SAP Datasphere to be able to connect to third party data sources powered by SAP Open Connectors.

Preparations in SAP BTP and SAP Open Connectors Account

1. Set up an SAP BTP account and enable the SAP Integration Suite service with the SAP Open Connectors capability.

   > **Note**
   >
   > You need to know your SAP BTP subaccount information (provider, region, environment, trial - yes/no) later to select the appropriate SAP BTP subaccount region in SAP Datasphere when integrating the SAP Open Connectors account in your space.

   For information about setting up an SAP BTP trial version with the SAP Integration Suite service, see Set Up Integration Suite Trial. To enable SAP Open Connectors, you need to activate the *Extend Non-SAP Connectivity* capability in the Integration Suite.
   For information about setting up SAP Integration Suite from a production SAP BTP account, see Initial Setup in the *SAP Integration Suite* documentation.
2. In your SAP Open Connectors account, create connector instances for the sources that you want to connect to SAP Datasphere.
   For more information about creating an instance, see Authenticate a Connector Instance (UI) in the *SAP Open Connectors* documentation.
   For more information about connector-specific setup and connector-specific properties required to create an instance, see Connectors Catalog in the *SAP Open Connectors* documentation. There, click the connector in question and then *<connector name> API Provider Setup* or *<connector name> Authenticate a Connector Instance*.
3. In your SAP Open Connectors account, record the following information which you will require later in SAP Datasphere:
   - Organization secret and user secret - required when integrating the SAP Open Connectors account in your space.
   - Name of the connector instance - required when selecting the instance in the connection creation wizard

Preparations in SAP Datasphere

1. In the side navigation area, click (*Connections*), select a space if necessary, click the *SAP Open Connectors* tab, and then click *Integrate your SAP Open Connectors Account* to open the *Integrate your SAP Open Connectors Account* dialog.
2. In the dialog, provide the following data:
   1. In the *SAP BTP Sub Account Region* field, select the appropriate entry according to your SAP BTP subaccount information (provider, region, environment, trial - yes/no).

2. Enter your SAP Open Connectors organisation and user secret.
3. Click *OK* to integrate your SAP Open Connectors account with SAP Datasphere.

Results

With connection type *Open Connectors* you can now create connections to the third-party data sources available as connector instances with your SAP Open Connectors account.

5.20 Prepare Connectivity to Oracle

To be able to successfully validate and use a connection to an Oracle database for remote tables or data flows, certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the OracleLogReaderAdapter.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- An administrator has downloadad and installed the required JDBC library in the `<DPAgent_root>/lib` folder before registering the adapter with SAP Datasphere.
- Required Permissions for Oracle Trigger-Based Replication in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality Installation and Configuration Guide*

Data Flows

Before you can use the connection for data flows, the following is required:

- To directly consume data in data flows, the Oracle database must be available on the public internet.
- A DW administrator has uploaded the required ODBC driver file to SAP Datasphere.
  To use encrypted communication (connection is configured to use SSL), additional files are required to be uploaded.
  For more information, see Upload Third-Party ODBC Drivers (Required for Data Flows) [page 132].
- A DW administrator has uploaded the server certificate to SAP Datasphere.
  For more information, see Manage Certificates for Connections [page 131].

5.21 Prepare Connectivity to Precog

To be able to successfully validate and use a connection to Precog for view building certain preparations have to be made.

Before you can use the connection, the following is required:

- In Precog, you have added the source for which you want to create the connection.
- In SAP Datasphere, you have added the necessary Precog IP addresses to the IP allowlist. For more information, see Add IP address to IP Allowlist [page 128].

> **Note**
>
> You can find and copy the relevant IP addresses in the final step of the connection creation wizard.

5.22 Prepare Connectivity to SAP ABAP Systems

To be able to successfully validate and use a connection to an SAP ABAP system for remote tables or data flows, certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the ABAPAdapter.
  For the *Language* setting in the connection properties to have an effect on the language shown in the Data Builder, Data Provisioning Agent version 2.0 SP 05 Patch 10 (2.5.1) or higher is required.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- The ABAP user specified in the credentials of the SAP ABAP connection needs to have a specific set of authorizations in the SAP ABAP system. For more information, see: Authorizations in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.
- To access and copy data from SAP BW objects such as InfoProviders or characteristics, the appropriate authorization objects like S_RS_ADSO or S_RS_IOBJA are required for the ABAP user. For more information, see Overview: Authorization Objects in the *SAP NetWeaver* documentation.
  If you want to access data from SAP BW Queries, make sure that the ABAP user has the required analysis authorizations to read the data and that characteristic 0TCAIPROV (InfoProvider) in the authorization includes **@Q**, which is the prefix for Queries as InfoProviders. For more information, see Defining Analysis Authorizations in the *SAP NetWeaver* documentation.
- If you want to stream ABAP tables for loading large amounts of data without running into memory issues, you need to configure suitable security privileges for successful registration on an SAP Gateway and you need to create an RFC destination of type TCP/IP in the ABAP source system. With the RFC destination you

register the Data Provisioning Agent as server program in the source system. For more information, see Prerequisites for ABAP RFC Streaming [page 120].

- To be able to use ABAP Dictionary tables from connections to a SAP BW/4HANA system for remote tables and creating views, please make sure that SAP note 2872997 has been applied to the system.

Data Flows

> **Note**
>
> The availability of the data flow feature depends on the used version and Support Package level of the ABAP-based SAP system (SAP S/4HANA or the DMIS addon in the source). Make sure your source systems meet the minimum versions mentioned below.We recommend to use the latest available version of SAP S/4HANA and the DMIS add-on where possible and have the latest SAP notes and TCI notes implemented in your systems.
>
> For more information about integrating ABAP-based SAP systems and about required and recommended versions in the source systems, see SAP Note 2890171.

Before you can use the connection for data flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to your on-premise source. In the Cloud Connector configuration, an administrator has made sure that access to the required resources is granted.
  For more information, see Configure Cloud Connector [page 122].
  See also: SAP Note 2835207 (*ABAP connection type for SAP Data Intelligence*)
- When connecting to SAP S/4HANA Cloud to access CDS views, an administrator has created a communication arrangement for the communication scenario `SAP_COM_0532` (*SAP Data Hub – ABAP CDS Pipeline Integration*) in the SAP S/4HANA Cloud system.
  For more information, see Integrating CDS Views Using ABAP CDS Pipeline in the *SAP S/4HANA Cloud* documentation.
- When connecting to SAP S/4HANA to access CDS views, source version 1909 FPS01 plus SAP Note 2873666 or higher versions are supported.
- When connecting to ABAP-based systems using SAP LT Replication Server as the source to access tables:
  Connectivity to SAP LT Replication Server is supported for:
  - Systems with Addon DMIS 2011 Support Package 19 or higher (supporting SAP ECC 6.00 or higher)
  - Systems with Addon DMIS 2018 Support Package 4 or higher (supporting SAP S/4HANA 1709 or higher)
  - SAP S/4HANA 2020 or higher

  Source systems connected to an SAP LT Replication Server are supported down to:
  - version 4.6C via Addon DMIS 2010
  - version 6.20 via Addon DMIS 2011

  In the ABAP-based system in which SAP LT Replication Server is installed, an administrator has created a configuration to specify the source system, the SAP LT Replication Server system and SAP Datasphere as target system.
  For more information, see Creating a Configuration in the *SAP Landscape Transformation Replication Server* documentation.

Replication Flows

> **Note**
>
> The availability of the replication flow feature depends on the used version and Support Package level of the ABAP-based SAP system (SAP S/4HANA or the DMIS addon in the source). Make sure your source systems meet the minimum versions mentioned below.We recommend to use the latest available version of SAP S/4HANA and the DMIS add-on where possible and have the latest SAP notes and TCI notes implemented in your systems.
>
> For more information about integrating ABAP-based SAP systems and about required and recommended versions in the source systems, see SAP Note 2890171.

Before you can use the connection for replication flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to your on-premise source. For more information, see Configure Cloud Connector [page 122].
- For table-based replication via SAP LT Replication Server, replication is supported for systems using:
  - DMIS 2018 addon Support Package 6 or higher
  - DMIS 2020 addon Support Package 3 or higher

  Source systems connected to an SAP LT Replication Server are supported down to:
  - version 4.6C via Addon DMIS 2010
  - version 6.20 via Addon DMIS 2011

  > **Note**
  >
  > Consider the version dependency when using a dedicated SAP LT Replication Server. For more information, see the *Installation Guide* in the SAP Landscape Transformation Replication Server documentation.
- ODP-based replication is supported for systems using:
  - DMIS 2011 addon Support Package 23 or higher
  - DMIS 2018 addon Support Package 8 or higher
  - DMIS 2020 addon Support Package 4 or higher

5.23 Prepare Connectivity to SAP BW

To be able to successfully validate and use a connection to SAP BW for remote tables or data flows, certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the ABAPAdapter.
  For the *Language* setting in the connection properties to have an effect on the language shown in the Data Builder, Data Provisioning Agent version 2.0 SP 05 Patch 10 (2.5.1) or higher is required.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- The ABAP user specified in the credentials of the SAP ABAP connection needs to have a specific set of authorizations in the SAP ABAP system. For more information, see: Authorizations in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.
- To access and copy data from SAP BW objects such as InfoProviders or characteristics, the appropriate authorization objects like S_RS_ADSO or S_RS_IOBJA are required for the ABAP user. For more information, see Overview: Authorization Objects in the *SAP NetWeaver* documentation.
  If you want to access data from SAP BW Queries, make sure that the ABAP user has the required analysis authorizations to read the data and that characteristic 0TCAIPROV (InfoProvider) in the authorization includes **@Q**, which is the prefix for Queries as InfoProviders. For more information, see Defining Analysis Authorizations in the *SAP NetWeaver* documentation.
- If you want to stream ABAP tables for loading large amounts of data without running into memory issues, you need to configure suitable security privileges for successful registration on an SAP Gateway and you need to create an RFC destination of type TCP/IP in the ABAP source system. With the RFC destination you register the Data Provisioning Agent as server program in the source system. For more information, see Prerequisites for ABAP RFC Streaming [page 120].
- To be able to use ABAP Dictionary tables from connections to a SAP BW/4HANA system for remote tables and creating views, please make sure that SAP note 2872997 has been applied to the system.

Data Flows

Before you can use the connection for data flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to your on-premise source. In the Cloud Connector configuration, an administrator has made sure that access to the required resources is granted.
  For more information, see Configure Cloud Connector [page 122].

5.24 Preparing SAP BW/4HANA Model Transfer Connectivity

Accessing SAP BW/4HANA meta data and importing models into SAP Datasphere with a SAP BW/4HANA Model Transfer connection requires two protocols (or endpoints): Http and SAP HANA Smart Data Integration based on the SAP HANA adapter.

For accessing SAP BW/4HANA, http is used to securely connect to the SAP BW/4HANA system via Cloud Connector, and SAP HANA SQL is used to connect to the SAP HANA database of SAP BW/4HANA via Data Provisioning Agent. Using Cloud Connector to make http requests to SAP BW/4HANA requires a live data connection of type tunnel to SAP BW/4HANA.

For information on supported SAP BW/4HANA source versions, see Supported Source Versions for SAP BW/4HANA Model Transfer Connections [page 153].

Before creating a connection for SAP BW/4HANA Model Transfer in SAP Datasphere, you need to prepare the following:

1. In SAP BW/4HANA, make sure that the following services are active in transaction code SICF:
   - BW InA - BW Information Access Services:
     - /sap/bw/ina/GetCatalog
     - /sap/bw/ina/GetResponse
     - /sap/bw/ina/GetServerInfo
     - /sap/bw/ina/ValueHelp
     - /sap/bw/ina/BatchProcessing
     - /sap/bw/ina/Logoff
   - /sap/bw4
2. In SAP BW/4HANA, activate OData service ESH_SEARCH_SRV in *Customizing* (transaction SPRO) under *SAP NetWeaver* > *Gateway* > *OData Channel* > *Administration* > *General Settings* > *Activate and Maintain Services*.
3. Install and configure Cloud Connector. For more information, see Configure Cloud Connector [page 122].
4. In the side navigation area of SAP Datasphere, click *System* > *Administration* > *Data Source Configuration* > *Live Data Sources* and switch on *Allow live data to leave my network*.

   > **Note**
   >
   > If your SAP Datasphere tenant was provisioned prior to version 2021.03, click *(Product Switch)* > *Analytics* > *System* > *Administration* > *Data Source Configuration* > *Live Data Sources*.

   For more information, Set Up Cloud Connector in SAP Datasphere [page 127].
5. In the side navigation area of SAP Datasphere, open *System* > *Configuration* > *Data Integration* > *Live Data Connections (Tunnel)* and create a live data connection of type tunnel to SAP BW/4HANA.

   > **Note**
   >
   > If your SAP Datasphere tenant was provisioned prior to version 2021.03, click *(Product Switch)* > *Analytics* > *(Connections)*.

   For more information, see Create Live Data Connection of Type Tunnel [page 151].

6. Install and configure a Data Provisioning Agent and register the SAP HANA adapter with SAP Datasphere:
    - Install the latest Data Provisioning Agent version on a local host or updat your agent to the latest version. For more information, see Install the Data Provisioning Agent [page 115].
    - In SAP Datasphere, add the external IPv4 address of the server on which your Data Provisioning Agent is running, or in case you are using a network firewall add the public proxy IP address to the IP allowlist. For more information, see Add IP address to IP Allowlist [page 128].
    - Connect the Data Provisioning Agent to SAP Datasphere. For more information, see Connect and Configure the Data Provisioning Agent [page 117].
    - In SAP Datasphere, register the SAP HANA adapter with SAP Datasphere. For more information, see Register Adapters with SAP Datasphere [page 119].

5.24.1 Create Live Data Connection of Type Tunnel

To securely connect and make http requests to SAP BW/4HANA, you need to connect via Cloud Connector. This requires that you create a live data connection of type tunnel to the SAP BW/4HANA system.

Prerequisites

See the prerequisites 1 to 4 in Preparing SAP BW/4HANA Model Transfer Connectivity [page 150].

Procedure

1. In the side navigation area, click (*System*) > (*Configuration*) > *Data Integration*.

    > **Note**
    >
    > If your SAP Datasphere tenant was provisioned prior to version 2021.03, click *(Product Switch)* > *Analytics* > *(Connections)* and continue with step 3.

2. In the *Live Data Connections (Tunnel)* section, click *Manage Live Data Connections*.

    The *Manage Live Data Connections* dialog appears.

3. On the *Connections* tab, click + *(Add Connection)*.

    The *Select a data source* dialog will appear.

4. Expand *Connect to Live Data* and select *SAP BW*.

    The *New BW Live Connection* dialog appears.

5. Enter a name and description for your connection. Note that the connection name cannot be changed later.
6. Set the *Connection Type* to *Tunnel*.

    By enabling tunneling, data from the connected source will always be transferred through the Cloud Connector.

7. Select the Location ID.
8. Add your SAP BW/4HANA host name, HTTPS port, and client.

   Use the virtual host name and virtual port that were configured in the Cloud Connector.
9. Optional: Choose a *Default Language* from the list.

   This language will always be used for this connection and cannot be changed by users without administrator privileges.

   > **Note**
   >
   > You must know which languages are installed on your SAP BW/4HANA system before adding a language code. If the language code you enter is invalid, SAP Datasphere will default to the language specified by your system metadata.
10. Under *Authentication Method* select *User Name and Password*.
11. Enter user name (case sensitive) and password of the technical user for the connection.

    The user needs the following authorizations:

    - Authorization object S_BW4_REST (authorization field: BW4_URI, value: /sap/bw4/v1/dwc*)
    - Authorization object SDDLVIEW (authorization field: DDLSRCNAME, value: RSDWC_SRCH_QV)
    - Read authorizations for SAP BW/4HANA metadata (Queries, CompositeProviders and their InfoProviders)
      Using authorizations for SAP BW/4HANA metadata, you can restrict a model transfer connection to a designated semantic SAP BW/4HANA area.
      For more information, see Overview: Authorization Objects in the SAP BW/4HANA documentation.
12. Select *Save this credential for all users on this system*.
13. Click *OK*.

    > **Note**
    >
    > While saving the connection, the system checks if it can access /sap/bc/ina/ services in SAP BW/4HANA.

Results

The connection is saved and now available for selection in the SAP Datasphere connection creation wizard for the SAP BW/4HANA Model Transfer connection.

5.24.2 Supported Source Versions for SAP BW/4HANA Model Transfer Connections

In order to create a connection of type *SAP BW/4HANA Model Transfer* , the SAP BW/4HANA system needs to have a specific version.

These versions of SAP BW/4HANA are supported:

- SAP BW/4HANA 2.0 SPS07 or higher
  - 2989654 BW/4 - Enable DWC "Import from Connection" for BW/4 Query - Revision 1
  - 2714624 Version Comparison False Result
  - 2754328 Disable creation of HTTP Security Sessions per request
  - 2840529 Sporadic HTTP 403 CSRF token validation errors
  - 2976147 Import of query views in the BW/4 hybrid scenario: No search results of BW back ends with SAP_BASIS Release 753
- SAP BW/4HANA 2.0 SPS01 to SPS06 after you have applied the following SAP Notes:
  - 2943200 TCI for BW4HANA 2.0 Hybrid
  - 2945277 BW/4 - Enable DWC "Import from Connection" for BW/4 Query
  - 2989654 BW/4 - Enable DWC "Import from Connection" for BW/4 Query - Revision 1
  - 2714624 Version Comparison False Result
  - 2754328 Disable creation of HTTP Security Sessions per request
  - 2840529 Sporadic HTTP 403 CSRF token validation errors
  - 2976147 Import of query views in the BW/4 hybrid scenario: No search results of BW back ends with SAP_BASIS Release 753

5.25 Prepare Connectivity to SAP ECC

To be able to successfully validate and use a connection to SAP ECC for remote tables or data flows, certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the ABAPAdapter.
  For the *Language* setting in the connection properties to have an effect on the language shown in the Data Builder, Data Provisioning Agent version 2.0 SP 05 Patch 10 (2.5.1) or higher is required.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- The ABAP user specified in the credentials of the SAP ABAP connection needs to have a specific set of authorizations in the SAP ABAP system. For more information, see: Authorizations in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.

- If you want to stream ABAP tables for loading large amounts of data without running into memory issues, you need to configure suitable security privileges for successful registration on an SAP Gateway and you need to create an RFC destination of type TCP/IP in the ABAP source system. With the RFC destination you register the Data Provisioning Agent as server program in the source system. For more information, see Prerequisites for ABAP RFC Streaming [page 120].

Data Flows

Before you can use the connection for data flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to your on-premise source. In the Cloud Connector configuration, an administrator has made sure that access to the required resources is granted.
  For more information, see Configure Cloud Connector [page 122].

5.26 Prepare Connectivity to SAP Fieldglass

To be able to successfully validate and use a connection to SAP Fieldglass for remote tables or data flows, certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the CloudDataIntegrationAdapter.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].

5.27 Prepare Connectivity to SAP HANA

To be able to successfully validate and use a connection to SAP HANA Cloud or SAP HANA (on-premise) for remote tables or data flows certain preparations have to be made.

SAP HANA Cloud

A DW administrator has uploaded the server certificate to SAP Datasphere.

For more information, see Manage Certificates for Connections [page 131].

SAP HANA on-premise

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- If you want to use SAP HANA Smart Data Integration:
  - An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the HanaAdapter.
    For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
  - If you use encrypted communication (see the *Security* properties in the connection creation wizard):
    An administrator has already correctly configured Data Provisioning Agent for SSL support.
    For more information, see Configure SSL for SAP HANA On-Premise [Manual Steps] in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.
- If you want to use SAP HANA Smart Data Access:
  - An administrator has installed and configured Cloud Connector to connect to your on-premise source.
    For more information, see Configure Cloud Connector [page 122].
  - An administrator has added the Cloud Connector IP address to the IP allowlist.
    For more information, see Add IP address to IP Allowlist [page 128].
  - If you use encrypted communication and the server certificate should be validated (see the *Security* properties in the connection creation wizard):
    A DW administrator has uploaded the server certificate to SAP Datasphere.
    For more information, see Manage Certificates for Connections [page 131].

Data Flows and Replication Flows

For SAP HANA (on-premise), before you can use the connection for data flows and replication flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to your on-premise source.
  For more information, see Configure Cloud Connector [page 122].

5.28 Prepare Connectivity to SAP Marketing Cloud

To be able to successfully validate and use a connection to SAP Marketing Cloud for remote tables or data flows, certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the CloudDataIntegrationAdapter.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- A communication arrangement has been created for communication scenario `SAP_COM_0531` in the source system.
  For more information, see Integrating CDI in the *SAP Marketing Cloud* documentation.

Data Flows

Before you can use the connection for data flows, the following is required:

- A communication arrangement has been created for communication scenario `SAP_COM_0531` in the source system.
  For more information, see Integrating CDI in the *SAP Marketing Cloud* documentation.

5.29 Prepare Connectivity to SAP SuccessFactors

To be able to successfully validate and use a connection to SAP SuccessFactors for remote tables or data flows certain preparations have to be made.

Before you can use the connection, the following is required:

- A DW administrator has uploaded the server certificate to SAP Datasphere.
  For more information, see Manage Certificates for Connections [page 131].
- When using OAuth 2.0 for authentication:
  - SAP Datasphere must be registered in SAP SuccessFactors.
    For more information, see Registering Your OAuth2 Client Application in the *SAP SuccessFactors platform* documentation.
  - A SAML assertion needs to be generated to be able to provide it when creating or editing the connection.
    For an overview of the available options to generate a SAML assertion, see Generating a SAML Assertion in the *SAP SuccessFactors platform* documentation.

- In SAP SuccessFactors IP restriction management, you have added the externally facing SAP HANA IP addresses and the outbound IP address for SAP Datasphere to the list of IP restrictions. IP restrictions are a specified list of IP addresses from which users can access your SAP SuccessFactors system.
  For more information, see:
  - IP Restrictions in the *SAP SuccessFactors platform* documentation
  - Finding SAP Datasphere IP addresses [page 130]

5.30 Prepare Connectivity to SAP S/4HANA Cloud

To be able to successfully validate and use a connection to SAP S/4HANA Cloud, certain preparations have to be made.

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the CloudDataIntegrationAdapter.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- A communication arrangement has been created for communication scenario `SAP_COM_0531` in the source system.
  For more information, see Integrating CDI in the *SAP S/4HANA Cloud* documentation.

Data Flows and Replication Flows

Before you can use the connection for data flows and replication flows, the following is required:

- A communication arrangement has been created for communication scenario `SAP_COM_0532` in the SAP S/4HANA Cloud system.
  For more information, see Integrating CDS Views Using SAP Datasphere in the *SAP S/4HANA Cloud* documentation.

Model Import

Before you can use the connection for model import, the following is required:

- A connection to an SAP HANA Smart Data Integration (SDI) Data Provisioning Agent with a registered CloudDataIntegrationAdapter.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].

- In the SAP S/4HANA Cloud system, communication arrangements have been created for the following communication scenarios:
  - `SAP_COM_0532`
    For more information, see Integrating CDS Views Using SAP Datasphere in the *SAP S/4HANA Cloud* documentation.
  - `SAP_COM_0531`
    For more information, see Integrating CDI in the *SAP S/4HANA Cloud* documentation.
  - `SAP_COM_0722`
    For more information, see Integrating SAP Data Warehouse Cloud in the *SAP S/4HANA Cloud* documentation.

5.31 Prepare Connectivity to SAP S/4HANA On-Premise

To be able to successfully validate and use a connection to SAP S/4HANA, certain preparations have to be made.

This topic contains the following sections:

Remote Tables

Before you can use the connection for creating views and accessing data via remote tables, the following is required:

- An administrator has connected an SAP HANA smart data integration Data Provisioning Agent to SAP Datasphere and registered the ABAPAdapter.
  For the *Language* setting in the connection properties to have an effect on the language shown in the Data Builder, Data Provisioning Agent version 2.0 SP 05 Patch 10 (2.5.1) or higher is required.
  For more information, see Preparing Data Provisioning Agent Connectivity [page 113].
- The ABAP user specified in the credentials of the SAP ABAP connection needs to have a specific set of authorizations in the SAP ABAP system. For more information, see: Authorizations in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.
- If you want to stream ABAP tables for loading large amounts of data without running into memory issues, you need to configure suitable security privileges for successful registration on an SAP Gateway and you need to create an RFC destination of type TCP/IP in the ABAP source system. With the RFC destination you register the Data Provisioning Agent as server program in the source system. For more information, see Prerequisites for ABAP RFC Streaming [page 120].

Data Flows

Before you can use the connection for data flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to your on-premise source. For more information, see Configure Cloud Connector [page 122].
- Supported source versions: SAP S/4HANA version 1909 FPS01 plus SAP Note 2873666 (*SAP Data Hub / Data Intelligence ABAP Integration - TCI note for SAP_ABA 1909 SP0/SP1*) or a higher SAP S/4HANA version

Replication Flows

Before you can use the connection for replication flows, the following is required:

- An administrator has installed and configured Cloud Connector to connect to your on-premise source. For more information, see Configure Cloud Connector [page 122].
- For CDS views, replication is supported for version 1909 and later. TCI note implementation is required for:
  - SAP S/4HANA 1909 with TCI Note 3105890 and central note 2830276.
  - SAP S/4HANA 2020 with TCI Note 3105880 and central note 2943599.
  - SAP S/4HANA 2021 with TCI Note 3115128 and central note 3085579.

  For information about which ABAP CDS Views are available for extraction, see CDS Views Enabled for Data Extraction in the *SAP S/4HANA* documentation.
- For table-based replication via SAP LT Replication Server, replication is supported for SAP S/4HANA 2022 and higher.
- ODP-based replication is supported for ODP context ODP_SAPI and ODP_BW, and requires ODP API version 2.

For more information about integrating ABAP-based SAP systems, see 2890171.

Model Import

Before you can use the connection to import entities, the following is required:

In SAP S/4HANA

- An administrator has followed the instructions from SAP Note 3081998 to properly set up the SAP S/4HANA system, which includes:
  1. SAP Note 3283282 has been implemented to provide the required infrastructure in the SAP S/4HANA system.
  2. The required corrections have been implemented and checks have been performed to make sure that SAP Note 3283282 and subsequent corrections have been applied properly and all required objects to provide the infrastructure are available and activated.
  3. Report `ESH_CSN_CDS_TO_CSN` has been run to prepare the CDS Views for the import.

- An administrator has created a technical user with the following authorizations:
  - Authorization object S_SERVICE - service authorizations for the Enterprise Search search service

Field	Value
SRV_NAME	EF608938F3EB18256CE851763C2952
SRV_TYPE	HT

  - Authorization object SDDLVIEW - Search access authorization for the search view CSN_EXPOSURE_CDS

Field	Value
DDLNAME	<leave empty - this field is not used>
DDLSRCNAME	CSN_EXPOSURE_CDS
ACTVT	03

  - Authorizations for remote table access via ODP
- An adminstrator has checked that the required InA services are active in transaction code SICF:
  - /sap/bw/ina/GetCatalog
  - /sap/bw/ina/GetResponse
  - /sap/bw/ina/GetServerInfo
  - /sap/bw/ina/ValueHelp
  - /sap/bw/ina/BatchProcessing
  - /sap/bw/ina/Logoff
- An administrator has activated OData service ESH_SEARCH_SRV in *Customizing* (transaction SPRO) under *SAP NetWeaver* > *Gateway* > *OData Channel* > *Administration* > *General Settings* > *Activate and Maintain Services*.

Cloud Connector

- An administrator has installed and configured Cloud Connector to connect to your on-premise source.
  For more information, see Configure Cloud Connector [page 122].

Data Provisioning Agent

- For the remote tables that will be created during the import, the respective prerequisites have to be met including a Data Provisioning Agent with the ABAPAdapter registered in SAP Datasphere.
  For more information, see Remote Tables [page 158].

In SAP Datasphere

- In *System* > *Administration* > *Data Source Configuration* > *Live Data Sources*, you have switched on *Allow live data to securely leave my network*.
  For more information, Set Up Cloud Connector in SAP Datasphere [page 127].
- In *System* > *Configuration* > *Data Integration* > *Live Data Connections (Tunnel)*, you have created a live data connection of type tunnel to SAP S/4HANA.
  For more information, see Create SAP S/4HANA Live Data Connection of Type Tunnel [page 161].

Supported Source Versions

- Supported source versions: SAP S/4HANA 1809 or higher

Related Information

SAP S/4HANA On-Premise Connections

5.31.1 Create SAP S/4HANA Live Data Connection of Type Tunnel

To securely connect to SAP S/4HANA on-premise when searching for ABAP CDS Views to be imported with the *Import Entities* wizard, you need to connect via Cloud Connector. This requires that you create a live data connection of type tunnel to the SAP S/4HANA system.

Prerequisites

See: Model Import [page 159]

Procedure

1. In the side navigation area, click (*System*) (*Configuration*) *Data Integration* .

   > **Note**
   >
   > If your SAP Datasphere tenant was provisioned prior to version 2021.03, click *(Product Switch)* *Analytics* *(Connections)* and continue with step 3.

2. In the *Live Data Connections (Tunnel)* section, click *Manage Live Data Connections*.

   The *Manage Live Data Connections* dialog appears.

3. On the *Connections* tab, click + *(Add Connection)*.

   The *Select a data source* dialog appears.

4. Expand *Connect to Live Data* and select *SAP S/4HANA*.

   The *New S/4HANA Live Connection* dialog appears.

5. Enter a name and description for your connection. Note that the connection name cannot be changed later.

6. Set the *Connection Type* to *Tunnel*.

7. Select the Location ID.
8. Add your SAP S/4HANA host name, HTTPS port, and client.

   Use the virtual host name and virtual port that were configured in the Cloud Connector.
9. Optional: Choose a *Default Language* from the list.

   This language will always be used for this connection and cannot be changed by users without administrator privileges.

   > **Note**
   >
   > You must know which languages are installed on your SAP S/4HANA system before adding a language code. If the language code you enter is invalid, SAP Datasphere will default to the language specified by your system metadata.
10. Under *Authentication Method* select *User Name and Password*.
11. Enter user name (case sensitive) and password of the technical user for the connection.
12. Select *Save this credential for all users on this system*.
13. Click *OK*.

Results

The connection is saved and now available for selection in the SAP Datasphere connection creation wizard for the *SAP S/4HANA On-Premise* connection.

6 Managing and Monitoring Connectivity for Data Integration

Monitor Data Provisioning Agent connectivity in SAP Datasphere, manage the impacts of agent changes in SAP Datasphere, and troubleshoot Data Provisioning Agent or Cloud Connector connectivity.

6.1 Monitoring Data Provisioning Agent in SAP Datasphere

For connected Data Provisioning Agents, you can proactively become aware of resource shortages on the agent instance and find more useful information.

In *Configuration* > *Data Integration* > *On-Premise Agents* choose the *Monitor* button to display the agents with the following:

- Information about free and used physical memory and swap memory on the Data Provisioning Agent server.
- Information about when the agent was connected the last time.
- Information about the overall number of connections that use the agent and the number of connections that actively use real-time replication, with active real-time replication meaning that the connection type supports real-time replication and for the connection at least one table is replicated via real-time replication.
  You can change to the *Connections* view to see the agents with a list of all connections they use and their real-time replication status. You can pause real-time replication for the connections of the while applying changes to the agent. For more information, see Pause Real-Time Replication for an Agent [page 167].

6.1.1 Monitoring Data Provisioning Agent Logs

Access the Data Provisioning Agent adapter framework log and the adapter framework trace log directly in SAP Datasphere.

With the integrated log access, you don't need to leave SAP Datasphere to monitor the agent and analyze agent issues. Accessing the log data happens via the Data Provisioning Agent File adapter which reads the log files and saves them into the database of SAP Datasphere.

The following logs are available:

Log File Name and Location on Data Provisioning Agent Server	Description
`<DPAgent_root>/log/framework_alert.trc`	Data Provisioning Agent adapter framework log. Use this file to monitor data provisioning agent statistics.

Log File Name and Location on Data Provisioning Agent Server	Description
`<DPAgent_root>/log/framework.trc`	Data Provisioning Agent adapter framework trace log. Use this file to trace and debug data provisioning agent issues.

You can review the logs in SAP Datasphere after log access has been enabled for the agent in question. We display the actual log files as well as up to ten archived log files that follow the naming convention `framework.trc.<x>` respectively `framework_alert.trc.<x>` with `<x>` being a number between one and ten.

Related Information

Enable Access to Data Provisioning Agent Logs [page 164]
Review Data Provisioning Agent Logs [page 165]

6.1.2 Enable Access to Data Provisioning Agent Logs

Enable accessing an agent's log files before you can view them in SAP Datasphere.

Prerequisites

A Data Provisioning Agent administrator has provided the necessary File adapter configuration with an access token that you need for enabling the log access in SAP Datasphere.

To configure the access token in the agent's secure storage, the administrator has performed the following steps in the agent configuration tool in command-line interactive mode:

1. At the command line, navigate to `<DPAgent_root>/bin`.
2. Start the agent configuration tool with the **`setSecureProperty`** parameter.
   - On Windows: `agentcli.bat --setSecureProperty`
   - On Linux, `./agentcli.sh --setSecureProperty`
3. Choose *Set FileAdapter Access Token* and specify the token.

For more information about the File adapter configuration, see File in the *Installation and Configuration Guide* of the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.

Procedure

1. From the ☰ main menu, open *Configuration* > *Data Integration*.

2. On the agent's tile, click *Edit*.
3. In the *Agent Settings* dialog, set *Enable Log Access* to *true*.
4. In the *FileAdapter Password* field that appears, enter the File adapter access token.
5. Click *Save* to activate the log access.

Results

The *Review Logs* entry in the menu of the agent's tile is enabled and the `framework_alert.trc` and `framework.trc` logs are written to the database of SAP Datasphere. You can now review the current and archived log files from the agent's tile.

6.1.3 Review Data Provisioning Agent Logs

Use the logs to monitor the agent and analyze issues with the agent.

Prerequisites

The logs are written to the database of SAP Datasphere. For more information, see Enable Access to Data Provisioning Agent Logs [page 164].

Procedure

1. From the ☰ main menu, open *Configuration* > *Data Integration*.
2. On the agent's tile, click *Review Logs*.

   The *Review Agent Logs* dialog initially shows 50 log entries. To load another chunks of 50 entries each, scroll down to the bottom of the dialog and use the *More* button.
3. To show the complete message for a log entry, click *More* in the *Message* column.
4. You have the following options to restrict the results in the display of the logs:
   - Search: In the *<agent name>* field, enter a search string and click 🔍 *(Search)* to search in the messages of the logs.
   - Filters: You can filter based on time, message type and log file name. When you've made your selection, click *Apply Filters*.

     > **Note**
     >
     > If your local time zone differs from the time zone used in the Data Provisioning Agent logs and you're applying a time-based filter, you might get other filter results than expected.

5. [optional] Export the logs as CSV file to your local system. Note that filters and search restrictions will be considered for the exported file.

6.1.4 Receive Notifications About Data Provisioning Agent Status Changes

For a selected SAP HANA Smart Data Integration Data Provisioning Agent, you can configure to get notified when the agent's status changes from connected to disconnected or the other way round.

Prerequisites

To run recurring scheduled tasks on your behalf, you need to authorize the job scheduling component of SAP Datasphere. In your profile settings under *Schedule Consent Settings*, you can give and revoke your consent to SAP Datasphere to run your scheduled tasks in the future. Note that when you don't give your consent or revoke your consent, tasks that you own won't be executed but will fail.

For more information, see Changing SAP Datasphere Settings.

Context

A recurring task will check for any status changes according to the configured frequency and send the notifications to the user who is the owner of the configuration. The initial owner is the user who created the configuration. Any user with the appropriate administration privileges can take over the ownership for this task if required, for example in case of vacation replacement or when the previous owner left the department or company.

Procedure

1. In the side navigation area, click ⓘ (*System*) > 🔧 (*Configuration*) > *Data Integration*.
2.
3. Go to the *On-Premise Agents* section and click ••• *(menu)* > *Configure Sending Notifications*.
4. If you haven't authorized SAP Datasphere yet to run your scheduled tasks for you, you will see a message at the top of the *Configure Sending Notifications* dialog asking for your consent. Give your consent.
5. Switch on the *Send Notifications* toggle.

   An additional field *Owner* appears that shows that you have been automatically assigned as the owner of the task.
6. Select the frequency in which the status of the Data Provisioning Agent should be checked.

7. Save your configuration.

   This will start the first status check. After the first check, the status check will be performed according to the defined frequency.

Results

If the status check finds any status change for the agent, a notification will be sent that you can find by clicking *(Notifications)* on the shell bar.

When you click on the notification, you'll get to the *On-Premise Agents* section in (*System*) > (*Configuration*) > *Data Integration* where you can start searching for the root cause in case the agent is disconnected.

Next Steps

If you need to take over the ownership and receive the notifications for an agent's status changes, go the the *Configure Sending Notifications* dialog as described above, click *Assign to Me* and save the configuration. From now on you will receive the notifications about any status changes for the agent. If you haven't done so yet, you need to provide your consent before you can take over the ownership.

6.2 Pause Real-Time Replication for an Agent

For a selected SAP HANA Smart Data Integration Data Provisioning Agent, you can pause real-time replication for the connections that use the agent while applying changes to it, such as configuration changes or applying patches. After you have finished your agent changes, you can restart real-time replication.

Context

If you need to perform maintenance activities in a source system, you can pause real-time replication for the corresponding connection. For more information, see Pause Real-Time Replication for a Connection.

Procedure

1. In SAP Datasphere, from the main menu, open *Configuration* > *Data Integration* > *On-Premise Agents*.

2. To show the Data Provisioning Agent tiles with a list of all connections they use, click the *Connections* button.

   The real-time replication status of a connection shown here, can be:

Real-Time Replication status	When do we show the status?
Active	The connection type supports real-time replication and for the connection at least one table is replicated via real-time replication (even if the status in the *Remote Table Monitor* is *Error*).
Inactive	The connection type supports real-time replication and for the connection currently there is no table replicating via real-time replication.
Paused	The connection type supports real-time replication and for the connection at least for one table real-time replication is paused.

3. To pause the agent's connections with replication status *Active* or *Inactive*, on the tile of the agent choose ••• *(menu)* and then ⦿ *Pause All Connections*.

   In the list of connections shown on the tile, the status for affected connections changes to *Paused*. You can also see the status change for the connections in the *Connections* application.

   In the *Remote Table Monitor* the status for affected tables changes to *Paused* and actions related to real-time replication are not available for these tables. Also, you cannot start real-time replication for any table of a paused connection.
4. You can now apply the changes to your Data Provisiong Agent.
5. Once you're finished with the changes, restart real-time replication for the agent. Choose ••• *(menu)* and then ⦿ *Restart All Connections*.

   The status in the list of connections shown on the tile, in the *Connections* application as well as in the *Remote Table Monitor* changes accordingly and you can again perform real-time related actions for the tables or start real-time replication.

6.3 Troubleshooting the Data Provisioning Agent (SAP HANA Smart Data Integration)

If you encounter problems with the Data Provisioning Agent, you can perform various checks and take actions to troubleshoot the problems.

The following sections provide information about checks, logs, and actions that you can take to troubleshoot problems with the Data Provisionning Agent:

- Initial Checks [page 169]
- Configuration Checks [page 169]
- Logs and Traces [page 170]
- Performance [page 171]

- Validating the Connection from the Server the Agent is Running to SAP Datasphere [page 171]
- Troubleshooting Connection Issues [page 174]
- Reviewing Data Provisioning Agent Logs [page 176]
- SAP Notes [page 176]
- Support Information [page 176]

> **Note**
>
> In the following sections, filepaths and screenshots are based on a Linux-based installation of the agent. If you have installed the agent on a Microsoft Windows server, the slashes "/" must be replaced by backslashes "\".

Initial Checks

A Data Provisioning Agent administrator can perform the following checks:

- **Firewall**
  For a successful connection, make sure that outbound connections from the Data Provisioning Agent to the target host and port, which is provided in the Data Provisioning Agent registration information in SAP Datasphere, are not blocked by your firewall.
- **Agent version**
  Make sure to always use the latest released version of the Data Provisioning Agent. For information on supported and available versions for the Data Provisioning Agent, see the SAP HANA Smart Data Integration Product Availability Matrix (PAM).
  Make sure that all agents that you want to connect to SAP Datasphere have the same latest version.
- **Java Installation**
  Check whether a Java installation is available by running the command `java -version`. If you receive a response like `java: command not found`, use the Java installation which is part of the agent installation. The Java executable can be found in folder `<DPAgent_root>/sapjvm/bin`.

Configuration Checks

The agent configuration is stored in the `<DPAgent_root>/dpagentconfig.ini` file in the agent installation root location (`<DPAgent_root>`). A Data Provisioning Agent administrator can double-check for the correct values (please do not maintain the parameters directly in the configuration file; the values are set with the command-line agent configuration tool):

dpagentconfig.ini file	Agent Settings in SAP Datasphere
agent.name=`<Agent Name>`	*Agent Name* (the name defined by the user who registered the agent in SAP Datasphere; the name is case sensitive)
hana.port=`<HANA Port>`	*HANA Port*

dpagentconfig.ini file	Agent Settings in SAP Datasphere
hana.onCloud=false	n/a
hana.useSSL=true	*HANA Use SSL*
hana.server=<HANA Server>	*HANA Server*
jdbc.enabled=true	*HANA via JDBC*
jdbc.host=<HANA Server>	*HANA Server*
jdbc.port=<HANA Port>	*HANA Port*
jdbc.encrypt=true	n/a

If you use a proxy server in your landscape, additionally check for the following parameters:

dpagentconfig.ini file
proxyType=http
jdbc.useProxy=true
jdbc.proxyHost=<your proxy host>
jdbc.proxyPort=<your proxy port>
jdbc.proxyHttp=true (true in case of http proxy, false in case of SOCKS proxy)
[if proxy authentication is required] jdbc.useProxyAuth=true
[if proxy authentication is required] jdbc.proxyUsername=<your proxy user name>
[if proxy authentication is required] jdbc.proxyPassword=<your proxy password>

For more information, see Agent Configuration Parameters in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.

Logs and Traces

To troubleshoot connection issues, a Data Provisioning Agent administrator can enable logging and JDBC tracing for the Data Provisioning Agent.

- **Agent Logs**
  Change the logging level to INFO (default), ALL, DEBUG, or TRACE according to your needs. For more informatiaon, see SAP Note 2496051 - *How to change "Logging Level" (Trace level) of a Data Provisioning Agent - SAP HANA Smart Data Integration*.
  The parameters for the logging level in the `<DPAgent_root>/dpagentconfig.ini` file are:
  - framework.log.level

- service.log.level

> **Note**
>
> Changing the level to DEBUG or ALL will generate a large amount of data. We therefore recommend to change the logging level to these values only for a short period of time while you are actively debugging and change it to a lower information level after you have finished debugging.

See also SAP Note 2461391 - *Where to find Data Provisioning Agent Log Files*

- **JDBC Trace**
  For information about activating JDBC tracing, see Trace a JDBC Connection in the *SAP HANA Service for SAP BTP in AWS and Google Cloud Regions* documentation.
  To set the trace level, execute the JDBC driver *.jar file from the `<DPAgent_root>/plugins` directory.

Performance

If you experience performance issues when replicating data via the Data Provisioning Agent, a Data Provisioning Agent administrator can consider increasing the agent memory as described in SAP Note 2737656 - *How to increase DP Agent memory*.

For general memory sizing recommendations for SAP HANA Smart Data Integration, see

- Data Provisioning Agent - Best Practices and Sizing Guide in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.
- SAP Note 2688382 - *SAP HANA Smart Data Integration Memory Sizing Guideline*

Validating the Connection from the Server the Agent is Running to SAP Datasphere

Ensure that your Data Provisioning Agent is connected to SAP HANA.

In SAP Datasphere

In (*System*) (*Configuration*) *Data Integration* *On-Premise Agents* a green bar and status information on the agent tile indicates if the agent is connected.

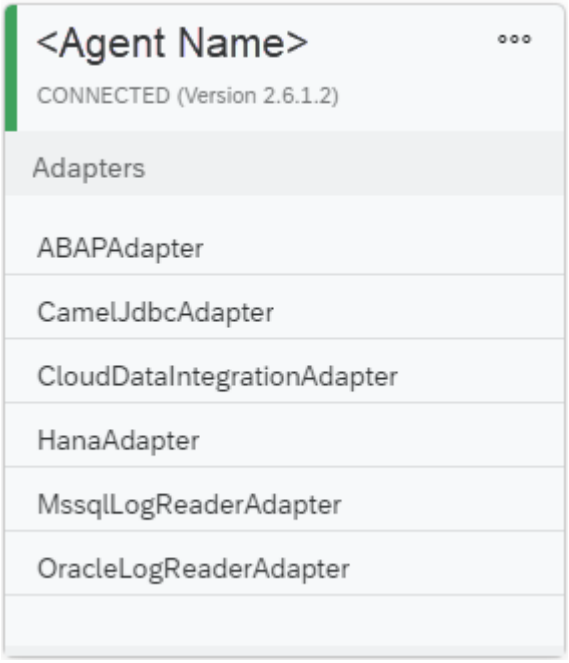

In *On-Premise Agents*, click *Refresh Agents* if the tile of a newly connected agent doesn't display the updated connection status.

> **Note**
>
> When you connect a new agent, it might take several minutes until it is connected.

Via Data Provisioning Agent Configuration Tool

1. Navigate to the command line and run `<DPAgent_root>/bin/agentcli.bat --configAgent`.
2. Choose *Agent Status* to check the connection status.
3. Make sure the output shows *Agent connected to HANA: Yes*.
4. If the output doesn't show that the agent is connected, it may show an error message. Resolve the error, and then select option *Start or Stop Agent*, and then option *Start Agent* to start the agent.

Via Trace File

The Data Provisioning Agent framework trace file `framework.trc` in the `<DPAgent_root>/log/` folder should contain entries indicating that the agent has been successfully connected.

```
2019-08-20 15:14:38,326 [INFO ] DPFramework | JDBCService.start [] - Starting JDBC service.
2019-08-20 15:14:38,326 [TRACE] DPFramework | JDBCConnector.read [] - Fetching messages...
2019-08-20 15:14:38,326 [INFO ] DPFramework | JDBCConnector.registerAgent [] - Registering
Agent: 78:69:REGISTER_AGENT:SDA:0:PAG_DPA12:null:null:null:0:0:REGISTER_AGENT:
2019-08-20 15:14:38,326 [DEBUG] DPFramework | JDBCConnector.writeMessage [] -
>>78:69:REGISTER_AGENT:SDA:0:PAG_DPA12:null:null:null:0:0:REGISTER_AGENT:
2019-08-20 15:14:38,904 [DEBUG] DPFramework | JDBCConnector.makeConnection [] - JDBC
connection (1) created.
2019-08-20 15:14:38,998 [DEBUG] DPFramework | JDBCConnector.writeMessage [] -
<<118:69:SUCCESS_MSG:SDA:28358:PAG_DPA12::::0:0:REGISTER_AGENT:
2019-08-20 15:14:38,998 [INFO ] DPFramework | JDBCConnector.registerAgent [] - Registration
complete.
2019-08-20 15:14:38,998 [INFO ] DPFramework | JDBCConnector.registerAgent [] -  Server Protocol:
SUPPORTS_2SP02_04_PROTOCOL
2019-08-20 15:14:38,998 [TRACE] DPFramework | JDBCConnector.getRegistrationToken [] - Token =
0000c134-f010-92f9-0000-000000000026
2019-08-20 15:14:38,998 [DEBUG] DPFramework | JDBCConnector.writeMessage [] -
>>106:105:GET_AGENT_MSG:SDA:0:PAG_DPA12:null:null:0000c134-f010-92f9-0000-
000000000026:0:0:GET_AGENT_MSG:
2019-08-20 15:14:39,373 [DEBUG] DPFramework | JDBCConnector.makeConnection [] - JDBC
connection (2) created.
```

Via Command Line

To validate the connection, you can directly use the JDBC driver jar file from the command line interface. You must ensure that you're using the same JDBC driver as used by the Data Provisioning Agent. The JDBC driver jar file (`com.sap.db.jdbc_*.jar`) is located in the `<DPAgent_root>/plugins` directory.

The pattern for the command line is:

```
java -jar <com.sap.db.jdbc_*.jar> -u <HANA User Name for Messaging Agent>,"<HANA
User Password for Messaging Agent>" -n <HANA Server>:<HANA Port> -o encrypt=true
```

Navigate to the `<DPAgent_root>/plugins/` directory and run one of the following commands by replacing the variables as needed and depending on your landscape:

- Without proxy:

  ```
  ../sapjvm/bin/java -jar <com.sap.db.jdbc_*.jar> -u <HANA User Name for
  Messaging Agent>,"<HANA User Password for Messaging Agent>" -n <HANA
  Server>:<HANA Port> -o encrypt=true
  ```

- With proxy:

  ```
  ../sapjvm/bin/java -jar <com.sap.db.jdbc_*.jar> -u <HANA User Name for
  Messaging Agent>,"<HANA User Password for Messaging Agent>" -n <HANA
  Server>:<HANA Port> -o encrypt=true -o proxyHostname=<your proxy host> -o
  proxyPort=<your proxy port> -o proxyHttp=true -o proxytype=http
  ```

- With proxy with authentication required:

  ```
  ../sapjvm/bin/java -jar <com.sap.db.jdbc_*.jar> -u <HANA User Name for
  Messaging Agent>,"<HANA User Password for Messaging Agent>" -n <HANA
  Server>:<HANA Port> -o encrypt=true -o proxyHostname=<your proxy host>
  ```

```
-o proxyPort=<your proxy port> -o proxyHttp=true -o proxytype=http -o
proxyUserName=<your proxy user name> -o proxyPassword="<your proxy password>"
```

If the connection works properly the statement should look like this:

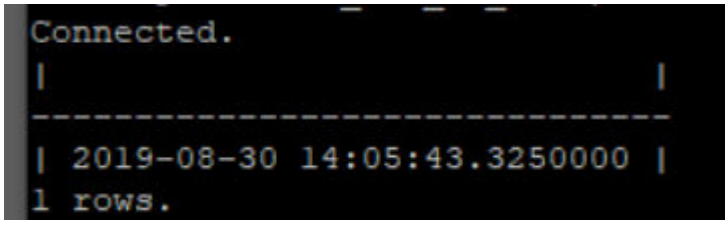

Troubleshooting Connection Issues

If you are unable to connect your Data Provisioning Agent to SAP Datasphere and have already validated the connection as described in the previous section, open the agent framework trace file `framework.trc` in the `<DPAgent_root>/log/` folder and check whether the output matches any of the following issues.

An entry is missing in the SAP Datasphere IP Allowlist

Example for an entry in the `framework.trc` file:

SQLException: SAP DBTech JDBC: Cannot connect to jdbc:sap://9e773134-7a64-44ee-93a1-4b34a06d4659.hana.prod-us10.hanacloud.ondemand.com:443 [Object is closed: com.sap.db.jdbc.SecureSession@9e773134-7a64-44ee-93a1-4b34a06d4659.hana.prod-us10.hanacloud.ondemand.com:443 ConnectionID:0 SessionID:0 on Socket[addr=9e773134-7a64-44ee-93a1-4b34a06d4659.hana.prod-us10.hanacloud.ondemand.com/<IP Address>,port=443,localport=53446]].

If you see this kind of error, it is most likely related to a missing entry in the IP Allowlist inSAP Datasphere.

Verify that the external (public) IPv4 address of the server where the agent is installed is in the IP allowlist. When using a proxy, the proxy's address needs to be included in IP allowlist as well.

For more information, see:

- Add IP address to IP Allowlist [page 128]
- SAP Note 2938870 - *Errors when connecting DP Agent with DWC*

Authentication failed

Example for an entry in the `framework.trc` file:

```
2019-08-20 14:51:03,508 [INFO ] DPFramework | JDBCService.start [] - Starting JDBC service.
2019-08-20 14:51:03,508 [TRACE] DPFramework | JDBCConnector.read [] - Fetching messages...
2019-08-20 14:51:03,508 [INFO ] DPFramework | JDBCConnector.registerAgent [] - Registering
Agent: 78:69:REGISTER_AGENT:SDA:0:PAG_DPA11:null:null:null:0:0:REGISTER_AGENT:
2019-08-20 14:51:03,508 [DEBUG] DPFramework | JDBCConnector.writeMessage [] -
>>78:69:REGISTER_AGENT:SDA:0:PAG_DPA11:null:null:null:0:0:REGISTER_AGENT:
2019-08-20 14:51:04,179 [ERROR] DPFramework | JDBCConnector.makeConnection [] - Failed
to connect to server (1 of 10 attempts).
2019-08-20 14:51:04,179 [ERROR] DPFramework | JDBCConnector.makeConnection [] -
com.sap.db.jdbc.exceptions.SQLInvalidAuthorizationSpecExceptionSapDB: [10]:
authentication failed
2019-08-20 14:51:04,179 [INFO ] DPFramework | JDBCConnector.makeConnection [] - Waiting for
30s to retry connection.
```

Authentication fails because of invalid *HANA User for Agent Messaging* credentials in the agent secure storage. To update the credentials, use the agent configuration tool and then restart the agent.

For more information, see Manage the HANA User for Agent Messaging Credentials in the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality* documentation.

Firewall/Proxy Issues

Example for an entry in the `framework.trc` file:

```
2019-08-20 14:08:41,433 [INFO ] DPFramework | JDBCService.start [] - Starting JDBC service.
2019-08-20 14:08:41,433 [TRACE] DPFramework | JDBCConnector.read [] - Fetching messages...
2019-08-20 14:08:41,433 [INFO ] DPFramework | JDBCConnector.registerAgent [] - Registering
Agent: 77:68:REGISTER_AGENT:SDA:0:PAG_DPA9:null:null:null:0:0:REGISTER_AGENT:
2019-08-20 14:08:41,433 [DEBUG] DPFramework | JDBCConnector.writeMessage [] -
>>77:68:REGISTER_AGENT:SDA:0:PAG_DPA9:null:null:null:0:0:REGISTER_AGENT:
2019-08-20 14:08:43,792 [ERROR] DPFramework | JDBCConnector.makeConnection [] - Failed
to connect to server (1 of 10 attempts).
2019-08-20 14:08:43,792 [ERROR] DPFramework | JDBCConnector.makeConnection [] -
com.sap.db.jdbc.exceptions.JDBCDriverException: SAP DBTech JDBC: Cannot connect to
jdbc:sap://<Host>:<Port> [Unknown host <Host>:<Port> [<Proxy>], -709].
2019-08-20 14:08:43,792 [INFO ] DPFramework | JDBCConnector.makeConnection [] - Waiting for
30s to retry connection.
```

This issue typically indicates that the JDBC driver is not capable of resolving the SAP HANA server URL to connect to theSAP Datasphere tenant and/or to establish a correct outbound call. Please check your firewall/proxy settings and make sure to enable outbound connections accordingly.

Encryption is missing: Only Secure Connections are Allowed

In case of missing encryption the log containts the following statement: "only secure connections are allowed".

```
SQLException: SAP DBTech JDBC: [4321]: only secure connections are allowed
```

When testing the connectivity directly with the JDBC driver, add the parameter `-o encrypt=true`.

Reviewing Data Provisioning Agent Logs

The logs are located in the `<DPAgent_root>/log` directory. For more information on the available log files, see SAP Note 2461391.

If the agent is connected, you can review the framework log (`framework_alert.trc`) and the framework trace log (`framework.trc`) directly in SAP Datasphere. For more information, see Monitoring Data Provisioning Agent Logs [page 163].

SAP Notes

SAP Note 2938870 - *Errors when connecting DP Agent with DWC*

SAP Note 2894588 - *IP Allowlist in SAP Datasphere*

SAP Note 2511196 - *What ports are used by Smart Data Integration*

SAP Note 2091095 - *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality*

SAP Note 2400022 - *FAQ: SAP HANA Smart Data Integration (SDI)*

SAP Note 2477204 - *FAQ: SAP HANA Services and Ports*

SAP Note 2688382 - *SAP HANA Smart Data Integration Memory Sizing Guideline*

Support Information

Support Component: SDI HAN-DP-SDI

Add and attach the following information:

- Version of the Data Provisioning Agent
- Framework trace log file (`framework.trc`)
- Data Provisioning Agent configuration file (`dpagentconfig.ini` file)

6.4 Troubleshooting Cloud Connector Related Issues

For information about troubleshooting Cloud Connector related issues when creating or using a connection in SAP Datasphere, see 3369433.

6.5 Troubleshooting SAP HANA Smart Data Access via Cloud Connector

These are some of the most common issues that can occur when you use the Cloud Connector to connect to on-premise remote sources via SAP HANA Smart Data Access.

1. The connectivity proxy is not enabled

The following error occurs if you try to connect to a remote source using the Cloud Connector, but the connectivity proxy hasn't been enabled:

```
[LIBODBCHDB SO][HDBODBC] Communication link failure;-10709 Connection failed
(RTE:[89001]
Cannot resolve host name '<connectivity_proxy_host>' rc=-2:
Name or service not known (<virtual_host>:<virtual_port>))
```

SAP Datasphere takes care of enabling the connectivity proxy. This might take a while.

2. The connectivity proxy is enabled but not fully ready to serve requests

The following error occurs if the connectivity proxy has been enabled but is not yet ready to be used:

```
[LIBODBCHDB SO][HDBODBC] Communication link failure;-10709 Connection failed
(RTE:[89006]
System call 'connect' failed, rc=111:
Connection refused {<connectivity_proxy_ip>:<connectivity_proxy_port>)}
{ClientPort:<client_port>} (<virtual_host>:<virtual_port>))
```

SAP Datasphere takes care of enabling the connectivity proxy. This might take a while.

3. The virtual host specified in the connection details includes an underscore

The following error occurs if you've used a virtual host name with an underscore, for example, hana_01:

```
[LIBODBCHDB SO][HDBODBC] General error;-10719 Connect failed (invalid SERVERNODE
'hana_01:<virtual_host>:<virtual_port>')
```

Virtual host names must not contain underscores.

4. The virtual host specified in the connection details is unreachable

The following error occurs if the specified virtual host cannot be reached:

```
[LIBODBCHDB SO][HDBODBC] Communication link failure;-10709 Connection failed
(RTE:[89132]
Proxy server connect: connection not allowed by ruleset
(<virtual_host>:<virtual_port>))
```

5. The selected location ID is invalid.

The following error occurs if an invalid location ID was specified in the *Data Source Configuration* of the SAP Datasphere *Administration*:

```
[LIBODBCHDB SO][HDBODBC] Communication link failure;-10709 Connection failed
(RTE:[89133]
Proxy server connect: Network unreachable (<virtual_host>:<virtual_port>))
```

6. The Cloud Connector's IP is missing or is incorrectly specified in the SAP Datasphere IP allowlist for trusted Cloud Connector IPs

The following error occurs when the Cloud Connector's IP is not included in the allowlist list:

```
[LIBODBCHDB SO][HDBODBC] Communication link failure;-10709 Connection failed
(RTE:[89133]
Proxy server connect: Network unreachable (<virtual_host>:<virtual_port>))
```

7. The Cloud Connector certificate has expired

The following error occurs when the subaccount certificate used in the Cloud Connector has expired:

```
[LIBODBCHDB SO][HDBODBC] Communication link failure;-10709 Connection failed
(RTE:[89133]
Proxy server connect: Network unreachable (<virtual_host>:<virtual_port>))
```

You can find the related logs in the `ljs_trace.log` file in the Cloud Connector. For example:

```
2021-07-29 04:50:42,131
+0200#ERROR#com.sap.core.connectivity.tunnel.client.notification.NotificationClie
nt#notification-client-277-3#
#Unable to handshake with notification server
connectivitynotification.cf.sap.hana.ondemand.com/<virtual_host>:<virtual_port>
javax.net.ssl.SSLException: Received fatal alert: certificate_expired
```

For information about renewing a subaccount certificate, see Update the Certificate for a Subaccount in the *SAP BTP Connectivity* documentation.

8. The on-premise backend system requires TCP SSL

The following error occurs if the on-premise backend system requires TCP SSL:

```
[LIBODBCHDB SO][HDBODBC] Communication link failure;-10709 Connection failed
(RTE:[89008]
Socket closed by peer (<virtual_host>:<virtual_port>))
```

Related Information

Troubleshooting Connection Issues with the Cloud Connector (SAP HANA Cloud, SAP HANA Database documentation)

7 Creating a Database User Group

Users with the DW Administrator role can create database user groups in SAP Datasphere to allow users to work in a sandboxed area in the underlying SAP HANA Cloud database, unattached to any space. These users can transfer an existing data warehouse implementation into the SAP Datasphere database or do any other work in SAP HANA Cloud and then make it available to one or more spaces as appropriate.

Context

When creating a database user group, an administrator is also created. This administrator can create other users, schemas, and roles using SAP Datasphere stored procedures. The administrator and their users can create data entities (DDL) and ingest data (DML) directly into their schemas and prepare them for consumption by spaces.

For detailed information about user groups, see User Groups in the *SAP HANA Cloud* documentation.

> **Note**
>
> Users with the DW Space Administrator role can create database users, which are associated with their space (see Integrating Data via Database Users/Open SQL Schemas).

Procedure

1. In the side navigation area, click (*System*) (*Configuration*) *Database Access* *Database User Groups* .
2. On the *Database User Group* page, click *Create*.
3. Enter a suffix for your database user group and click *Create*.

   The group is created and the connection details and administrator credentials are displayed.

   If you want to work with the SAP HANA database explorer, you will need to enter your password to grant the explorer access to the database user group schema. When connecting to SAP HANA Cloud with other tools, users will need the following properties:
   - Database Group Administrator (name and password)
   - Host Name
   - Port
4. Click *Close* to close the dialog.

7.1 Create Users, Schemas, and Roles in a Database User Group

A database user group administrator can create users, schemas, and roles to organise and staff their group. Creating schemas and roles and granting, revoking, and dropping roles require the use of SAP Datasphere stored procedures.

This topic contains the following sections:

Log In With Your Database User Group Administrator

To connect to SAP HANA Cloud with the administrator, select your newly created user group in the list, and click *Open Database Explorer*, enter the password when requested, and click *OK*.

The SAP HANA database explorer opens with your database user group at the top level. You can now use the SQL editor to create users, roles and schemas.

You can review your privileges with the following statement:

```
select * from effective_privileges where user_name = current_user;
```

Create a User

You can create a user in your user group with the following statement:

```
CREATE USER <user_name> PASSWORD <pwd> SET USERGROUP <DBgroup_name>
```

> **Note**
>
> To avoid possible conflicts, we recommend that you use the prefix `DWCDBGROUP#<DBgroup_name>#` when naming users, schemas, and roles in your group (see Rules for Technical Names [page 106]).

In our example, we create a new user, `DWCDBGROUP#DWMIGRATE#BOB`, in our `DWMIGRATE` group:

```
CREATE USER DWCDBGROUP#DWMIGRATE#BOB password “Welcome1” set usergroup
“DWCDBGROUP#DWMIGRATE”;
```

Create a Schema

You can create a schema in your database user group by using the following SAP Datasphere stored procedure:

```
CALL "DWC_GLOBAL"."CREATE_USERGROUP_SCHEMA"
(
    SCHEMA_NAME => '<schema_name>',
    OWNER_NAME => '<user_name>'
                   );
```

> ### Note
>
> To avoid possible conflicts, we recommend that you use the prefix `DWCDBGROUP#<DBgroup_name>#` when naming users, schemas, and roles in your group (see Rules for Technical Names [page 106]).

The owner of the new schema must be a user of the database user group. If the owner name is set to null, then the database user group administrator is set as the owner.

In our example, we create a new schema, `DWCDBGROUP#DWMIGRATE#STAGING`, and set `BOB` as the owner:

```
CALL "DWC_GLOBAL"."CREATE_USERGROUP_SCHEMA"
(
    SCHEMA_NAME => 'DWCDBGROUP#DWMIGRATE#STAGING,
    OWNER_NAME => 'DWCDBGROUP#DWMIGRATE#BOB'
);
```

Create a Role

You can create a role in your database user group by using the following SAP Datasphere stored procedure:

```
CALL "DWC_GLOBAL"."CREATE_USERGROUP_ROLE"
(
    ROLE_SCHEMA_NAME => '<schema_name>',
    ROLE_NAME => '<role_name>'
);
```

> ### Note
>
> To avoid possible conflicts, we recommend that you use the prefix `DWCDBGROUP#<DBgroup_name>#` when naming users, schemas, and roles in your group (see Rules for Technical Names [page 106]).

Once the role is created, you can grant it to a user or to another role, revoke it, and drop it.

In our example, we create a new role, `DWCDBGROUP#DWMIGRATE#DWINTEGRATOR` in the schema `STAGING`:

```
CALL "DWC_GLOBAL"."CREATE_USERGROUP_ROLE"
(
    ROLE_SCHEMA_NAME => 'DWCDBGROUP#DWMIGRATE#STAGING',
    ROLE_NAME => 'DWCDBGROUP#DWMIGRATE#DWINTEGRATOR'
);
```

Grant a Role to a User or to Another Role

You can grant a role to a user or to another role in your database user group by using the following SAP Datasphere stored procedure:

```
CALL "DWC_GLOBAL"."GRANT_USERGROUP_ROLE"
(
    ROLE_SCHEMA_NAME => '<schema_name>',
    ROLE_NAME => '<role_name>',
    GRANTEE => '<user_name>',
    GRANTEE_ROLE_NAME => NULL,
    WITH_ADMIN_OPTION => FALSE
);
```

The role schema, grantee, and grantee role must all be in the same database user group.

In our example, we grant the `DWCDBGROUP#DWMIGRATE#DWINTEGRATOR` role to our user, `BOB`:

```
CALL "DWC_GLOBAL"."GRANT_USERGROUP_ROLE"
(
    ROLE_SCHEMA_NAME => 'DWCDBGROUP#DWMIGRATE#STAGING',
    ROLE_NAME => 'DWCDBGROUP#DWMIGRATE#DWINTEGRATOR',
    GRANTEE => 'DWCDBGROUP#DWMIGRATE#BOB',
    GRANTEE_ROLE_NAME => NULL,
    WITH_ADMIN_OPTION => FALSE
);
```

Revoke a Role

You can revoke a role from a user by using the following SAP Datasphere stored procedure:

```
CALL "DWC_GLOBAL"."REVOKE_USERGROUP_ROLE"
(
    ROLE_SCHEMA_NAME => '<schema_name>',
    ROLE_NAME => '<role_name>',
    GRANTEE => '<user_name>',
    GRANTEE_ROLE_NAME => NULL
);
```

In our example, we revoke the `DWCDBGROUP#DWMIGRATE#DWINTEGRATOR` role from `BOB`:

```
CALL "DWC_GLOBAL"."REVOKE_USERGROUP_ROLE"
(
    ROLE_SCHEMA_NAME => 'DWCDBGROUP#DWMIGRATE#STAGING',
    ROLE_NAME => 'DWCDBGROUP#DWINTEGRATOR',
    GRANTEE => 'DWCDBGROUP#DWMIGRATE#BOB',
    GRANTEE_ROLE_NAME => NULL
);
```

Drop a Role

You can drop a role by using the following SAP Datasphere stored procedure:

```
CALL "DWC_GLOBAL"."DROP_USERGROUP_ROLE"
(
ROLE_SCHEMA_NAME => '<schema_name>',
ROLE_NAME => '<role_name>'
        );
```

In our example, we drop the `DWCDBGROUP#DWMIGRATE#DWINTEGRATOR` role:

```
CALL "DWC_GLOBAL"."DROP_USERGROUP_ROLE"
(
ROLE_SCHEMA_NAME => 'DWCDBGROUP#DWMIGRATE#STAGING',
ROLE_NAME => 'DWCDBGROUP#DWMIGRATE#DWINTEGRATOR'
        );
```

7.2 Allow a Space to Read From the Database User Group Schema

By default, no SAP Datasphere space can access the database user group schema. To grant a space read privileges from the database user group schema, use the GRANT_PRIVILEGE_TO_SPACE stored procedure.

Prerequisites

Only the administrator of a database user group has the privilege to run the stored procedure "DWC_GLOBAL"."GRANT_PRIVILEGE_TO_SPACE".

Context

You can grant read privileges by running an SAP Datasphere specific stored procedure in the SQL console in the SAP HANA Database Explorer.

Procedure

1. From the side navigation area, go to ⓘ *(System)* → 🔧 (*Configuration*) → *Database Access* → *Database User Groups*.
2. Select the database user group and click *Open Database Explorer*.

3. In the SQL console in SAP HANA Database Explorer, call the stored procedure to grant the 'SELECT' privilege to a space using the following syntax:

```
CALL "DWC_GLOBAL"."GRANT_PRIVILEGE_TO_SPACE" (
    OPERATION => <operation>,
    PRIVILEGE => <privilege>,
    SCHEMA_NAME => <schema name>,
    OBJECT_NAME => <object name>,
SPACE_ID => <space ID>);
```

Parameters are set as follows:

Parameter	Values	Description
operation	• 'GRANT' • 'REVOKE'	[required] Enter 'GRANT' to give the read privileges, or 'REVOKE' to remove the read privileges to the space.
privilege	'SELECT'	[required] Enter the read privilege that you want to grant (or revoke) to the space.
schema_name	'[name of database user group schema]'	[required] Enter the name of the schema you want the space to be able to read from.
object_name	• ' ' • null • '[name of the objet]'	[required] You can grant the read privileges, either at the schema level or at the object level. • At the schema level (all objets in the schema): enter null or ' '. • At the object level: enter a valid table name.
space_id	'[ID of the space]'	[required] Enter the ID of the space you are granting the read privileges to.

To grant read access to all objects (tables) in the schema:

```
CALL "DWC_GLOBAL"."GRANT_PRIVILEGE_TO_SPACE" (
    OPERATION => 'GRANT',
    PRIVILEGE => 'SELECT',
    SCHEMA_NAME => 'SALE#ETL',
    OBJECT_NAME => '',
    SPACE_ID => 'SALES');
```

To grant read access to the table MY_TABLE:

```
CALL "DWC_GLOBAL"."GRANT_PRIVILEGE_TO_SPACE" (
    OPERATION => 'GRANT',
    PRIVILEGE => 'SELECT',
    SCHEMA_NAME => 'SALE#ETL',
    OBJECT_NAME => 'MY_TABLE',
    SPACE_ID => 'SALES');
```

4. Run the query by clicking ⓡ *(Run)* or press F8.

Results

If the run is successful, you receive a confirmation message in the *Result* pane. You can then open the *Data Builder*, create a data flow, and select the tables as sources.

7.3 Allow a Space to Write to the Database User Group Schema

To grant a space write privileges in the database user group schema, use the GRANT_PRIVILEGE_TO_SPACE stored procedure. Once this is done, data flows running in the space can select tables in the schema as targets and write data to them.

Prerequisites

Only the administrator of a database user group has the privilege to run the stored procedure "DWC_GLOBAL"."GRANT_PRIVILEGE_TO_SPACE".

Context

You can grant write privileges by running an SAP Datasphere specific stored procedure in the SQL console in the SAP HANA Database Explorer.

Procedure

1. From the side navigation area, go to ⓘ *(System)* → 🔧 (*Configuration*) → *Database Access* → *Database User Groups*.
2. Select the database user group and click *Open Database Explorer*.
3. In the SQL console in SAP HANA Database Explorer, call the stored procedure to grant the 'INSERT', 'UPDATE', or 'DELETE' privilege to a space using the following syntax:

```
CALL "DWC_GLOBAL"."GRANT_PRIVILEGE_TO_SPACE" (
    OPERATION => <operation>,
    PRIVILEGE => <privilege>,
    SCHEMA_NAME => <schema name>,
    OBJECT_NAME => <object name>,
SPACE_ID => <space ID>);
```

Parameters are set as follows:

Parameter	Values	Description
operation	• 'GRANT' • 'REVOKE'	[required] Enter 'GRANT' to give the write privileges, or 'REVOKE' to remove the write privileges to the space.
privilege	• 'INSERT" • 'UPDATE' • 'DELETE'	[required] Enter the write privilege that you want to grant (or revoke) to the space. **Note** You can grant one privilege at a time.
schema_name	'[name of database user group schema]'	[required] Enter the name of the schema you want the space to be able to write from.
object_name	• ' ' • null • '[name of the objet]'	[required] You can grant the write privileges, either at the schema level or at the object level. • At the schema level (all objets in the schema): enter null or ' '. • At the object level: enter a valid table name.
space_id	'[ID of the space]'	[required] Enter the ID of the space you are granting the write privileges to.

To grant update write access to all objects (tables) in the schema:

```
CALL "DWC_GLOBAL"."GRANT_PRIVILEGE_TO_SPACE" (
    OPERATION => 'GRANT',
    PRIVILEGE => 'UPDATE',
    SCHEMA_NAME => 'SALE#ETL',
    OBJECT_NAME => '',
    SPACE_ID => 'SALES');
```

To grant update write access to the table MY_TABLE:

```
CALL "DWC_GLOBAL"."GRANT_PRIVILEGE_TO_SPACE" (
    OPERATION => 'GRANT',
    PRIVILEGE => 'UPDATE',
    SCHEMA_NAME => 'SALE#ETL',
    OBJECT_NAME => 'MY_TABLE',
    SPACE_ID => 'SALES');
```

4. Run the query by clicking ⓡ *(Run)* or press F8.

Results

If the run is successful, you receive a confirmation message in the *Result* pane. You can then open the *Data Builder*, create a data flow, and select the tables as targets.

8 Monitoring SAP Datasphere

Administrators have access to various monitoring logs and views, and can create database analysis users, if necessary, to help troubleshoot issues.

This topic contains the following sections:

Click (*System Monitor*) to access the main monitoring tool. The *System Monitor* allows to monitor the performance of your system and identify storage, task, out-of-memory, and other issues across all spaces.

For example, you can see all the errors (such as failed tasks and out-of-memory errors) that occurred yesterday or the top five statements with the highest peak memory consumption.

> **Note**
>
> For optimal performance, it is recommended that you consider staggering the scheduled run time of tasks such as data flows and task chains that may contain these tasks. There is a limit on how many tasks can be started at the same time. If you come close to this limit, scheduled task runs may be delayed and, if you go beyond the limit, some scheduled task runs might even be skipped.

Monitoring Out-of-Memory Errors and Other Statement-Related Information

You can monitor out-of-memory errors and other information that are related to SAP HANA database SQL statements, depending on what you've specified in *(Configuration)* → *Monitoring*:

- If *Enable Expensive Statement Tracing* is not selected, then in *System Monitor* > *Dashboard* , you cannot see the widgets about out-of-memory errors and about other information related to statements. For example, you cannot see the widgets: *Out-of-Memory Errors*, *Top 5 Statements by Processing Memory Consumption*.
- If *Enable Expensive Statement Tracing* is not selected and none of the threshold options is selected, then in the tables of *System Monitor* > *Logs* , you cannot see any information about out-of-memory errors and about other information related to statements. For example, no information is displayed in the columns *Peak Memory (MiB)* and *Peak CPU (ms)*.

A maximum of 30.000 records by default (a number that you can change) are stored for expensive statements. It means that the thresholds also determine how long you are able to see the statements. If the thresholds are very low, many statements are traced in a short timeframe.

For this reason, we recommend the following:

- *Threshold Duration* - Specify 5 seconds.
- *Threshold Memory* - Specify 1.000 MB.
- *Threshold CPU Time* - Do not select this option.

You can change the maximum number of records to trace more data in the *System Monitor*. For example, if about 5 days are traced in the expensive statement tables and you don't want to change the thresholds, you can double the number of records in *In-Memory Tracing Records* so that about 10 days are traced.

For more information on enabling and configuring expensive statement tracing, see Analyze Monitoring Data in a Space [page 197].

Dashboard

The out-of-memory widgets and top-memory consumption widgets help you to identify issues. You can see detailed information about the issues by clicking *View Logs*, which takes you to the *Logs* tab. For example, you can find out if tasks have to be scheduled at another time so that high-memory consuming tasks do not run at the same time. If single tasks consume too much memory, some additional views may need to be persisted or the view partitioning may need to be used to lower the memory consumption.

The following information is available in the *Dashboard* tab:

- *Disk Used by Spaces for Storage* - Shows the total amount of disk storage used in all spaces, out of the total amount of disk storage assigned to all spaces.
  In *Disk Storage Used*, you can see a breakdown between:
  - *Data in Spaces*: All data that is stored in spaces.
  - *Audit Log Data*: Data related to audit logs (see Audit Logging).

    > **Note**
    >
    > Audit logs can grow quickly and consume a great deal of disk storage (see Delete Audit Logs [page 206]).

  - *Other Data*: Includes data stored in database user group schemas (see Creating a Database User Group [page 180]) and SAP HANA data (such as statistics schemas).
  - *Administrative Data*: Data used to administer the tenant and all spaces (such as space quota, space version). Includes all information stored in the central schemas (DWC_GLOBAL, DWC_GLOBAL_LOG, DWC_TENANT_OWNER).
- *Memory Used by Spaces for Storage* - Shows the total amount of memory storage used in all spaces, out of the total amount of memory storage assigned to all spaces.
- *Disk Assigned to Spaces for Storage* - Shows the total amount of disk storage assigned to all spaces.
- *Memory Assigned to Spaces for Storage* - Shows the total amount of memory storage assigned to all spaces.

For each of the key indicator widgets listed below, you can see detailed information by clicking the link *View Logs*, which takes you to the *Logs* tab.

Widget	Description
Failed Tasks *Last 24 Hours*	Shows the number of tasks that have failed in the last 24 hours with a trend icon (up or down arrow) indicating if there are more or less failed tasks than the day before.
Out-of-Memory Errors *Last 24 Hours*	Shows the number of out-of-memory errors that have occurred in tasks and statements in the last 24 hours.
Top 5 Out-of-Memory Errors (Workload Class) by Space *Last 7 Days*	Shows the schemas in which out-of-memory errors have occurred in the last 7 days because the statement limits have been exceeded. You can set the statement limits for a space in the *Workload Management* area of a space.
Out-of-Memory Errors (MDS Requests) *Last 7 Days*	Shows the out-of-memory errors that are related to SAP HANA multi-dimensional services (MDS) requests, which is used for example for SAP Analytics Cloud consumption.
Failed Tasks *Last 7 Days*	Shows the number of failed tasks by day for the last 7 days.
Out-of-Memory Errors *Last 7 Days*	Shows the number of out-of-memory errors that have occurred in tasks and statements, by day for the last 7 days.
Top 5 Tasks by Run Duration *Last 24 Hours*	Shows the 5 tasks whose run duration time was the longest in the last 24 hours.
Top 5 Tasks by Run Duration *Last 48 Hours*	Shows the 5 tasks whose run duration time was the longest in the last 48 hours.
Top 5 Tasks by Processing Memory Consumption *Last 24 Hours*	Shows the 5 tasks whose processing memory consumption was the highest in the last 24 hours.
Top 5 Tasks by Processing Memory Consumption *Last 48 Hours*	Shows the 5 tasks whose processing memory consumption was the highest in the last 48 hours.
Top 5 Statements by Processing Memory Consumption *Last 24 Hours*	Shows the 5 statements whose processing memory consumption was the highest in the last 24 hours.
Top 5 Statements by Processing Memory Consumption *Last 48 Hours*	Shows the 5 statements whose processing memory consumption was the highest in the last 48 hours.
Top 5 MDS Requests by Processing Memory Consumption *Last 24 Hours*	Shows the 5 SAP HANA multi-dimensional services (MDS) requests (used for example in SAP Analytics Cloud consumption), whose processing memory consumption is the highest.

Widget	Description
Admission Control Rejection Events *Last 24 Hours*	Shows the number of statements that have been rejected in the last 24 hours because they've exceeded the threshold percentage of CPU usage. A trend icon (up or down arrow) indicates if there are more or less rejected statements than the day before.
Admission Control Queuing Events *Last 24 Hours*	Shows the number of statements that have been queued in the last 24 hours because they've exceeded the threshold percentage of CPU usage. A trend icon (up or down arrow) indicates if there are more or less queued statements than the day before.
Admission Control Rejection Events *Last 7 Days*	Shows the number of statements that have been rejected in the last 7 days because they've exceeded the threshold percentage of CPU usage.
Admission Control Queuing Events *Last 7 Days*	Shows the number of statements that have been queued in the last 7 days because they've exceeded the threshold percentage of CPU usage.
Top 5 Admission Control Rejection Events by Space *Last 7 Days*	Shows the 5 spaces with the highest number of rejected statements in the last 24 hours. **Note** A space that has been deleted is prefixed with an asterisk character.
Top 5 Admission Control Queuing Events by Space *Last 7 Days*	Shows the 5 spaces with the highest number of queued statements in the last 24 hours. **Note** A space that has been deleted is prefixed with an asterisk character.

Note

- To investigate why statements are being queued or rejected, you can click *Open SAP HANA Cockpit* in the widgets dedicated to admission contol. If you've created a database analysis user, you're connected to the SAP HANA Cockpit without entering your credentials (see Create a Database Analysis User to Debug Database Issues [page 208].
- For more information about admission control thresholds, see Set a Priority and Statement Limits for a Space [page 103].

Tasks

In *Logs* > *Tasks*, the table shows the following information:

Property	Description
Start Time	Shows at what time (date and hour) the task has started to run.
Duration (sec)	Shows how many seconds the task has run.

Property	Description
Object Type	Shows the type of object that was run in the task. For example: view, remote table, data flow.
Activity	Shows the action that was performed on the object. For example: persist, replicate, execute. You can click on the activity name, which takes you to the *Data Integration Monitor*.
Space Name	Shows the name of the space in which the task is run.
Object Name	Shows the name of the object. You can click on the object name, which opens the object in the *Data Builder*.
SAP HANA Peak Memory	Shows the maximum amount of memory (in MiB) the task has used during the runtime in SAP HANA. **Note** You can see this information: • If the option *Enable Expensive Statement Tracing* is enabled and if the task exceeds the thresholds specified in *(Configuration)* → *Monitoring*. • And if the task is run for these objects (and activities): views (persist, remove_persisted_data), remote tables (replicate, enable_realtime), data flows (execute) and intelligent lookup (execute, delete_data). Otherwise, no number is displayed.
SAP HANA CPU Time	Shows the amount of CPU time (in ms) the task has used in SAP HANA. **Note** You can see this information: • If the option *Enable Expensive Statement Tracing* is enabled and if the task exceeds the thresholds specified in *(Configuration)* → *Monitoring*. • And if the task is run for these objects (and activities): views (persist, remove_persisted_data), remote tables (replicate, enable_realtime), data flows (execute) and intelligent lookup (execute, delete_data). Otherwise, no number is displayed. **Note** The CPU time indicates how much time is used for all threads. It means that if the CPU time is significantly higher than the duration of the statement, then many threads are used. If many threads are used for a long time, no other tasks should be scheduled at that point in time, or resource bottlenecks may occur and tasks may even be canceled.
Records	Shows the number of records of the target table after the task has finished running. **Note** You can see this information only if the task is run for these objects (and activities): views (persist), remote tables (replicate, enable_realtime), data flows (execute) and intelligent lookup (execute, delete_data). Otherwise, no number is displayed.

Property	Description
SAP HANA Used Memory	Shows the amount of memory (in MiB) that is used by the target table in SAP HANA after the task has finished running.
SAP HANA Used Disk	Shows the amount of disk space (in MiB) that is used by the target table in SAP HANA after the task has finished running.
Status	Shows the status of the task: completed, failed, running.
Substatus	For tasks with the status "failed", shows the substatus and a message describing the cause of failure. For more information about failed task substatuses, see Understanding Statuses and Substatuses.
User	Shows the user who has run the task.
Target Table	Shows the SAP HANA database technical name of the target table.
Statements	Shows a link you can click to view all the statements of the task in the *Statements* table. **Note** • You can see this information if the option *Enable Expensive Statement Tracing* is enabled in *(Configuration)* → *Monitoring*. • However, as statements are traced for a limited period, you may not be able to see the statements used in the task.
Out-of-Memory	Shows if the task has an out-of-memory error ("Yes" is then displayed) or not ("No" is then displayed).
Task Log ID	Shows the identifier of the run task.

Note

Data on tasks are kept for the time specified in *(Configuration)* → *Tasks*.

Statements

In *Logs* > *Statements*, the table shows the following information, depending on what you've specified in *(Configuration)* → *Monitoring*:

- If the option *Enable Expensive Statement Tracing* is disabled, then the *Statements* tab is disabled.
- If the option *Enable Expensive Statement Tracing* is enabled, you can see all the database statements that exceed the specified thresholds.

Property	Description
Start Time	Shows at what time (date and hour) the statement has started to run.
Duration (ms)	Shows how many milliseconds the statement has run.

Property	Description

Object Type	• Shows the type of object that was run in the statement (for example: view, remote table, data flow).
	• Or shows the area where the statement was run:
	  • MDS - this is an SAP HANA multi-dimensional services (MDS) statement, which is caused for example by stories when SAP Analytics Cloud queries SAP Datasphere.
	  • Data Flow - the statement was run by a data flow.
	  • Analysis - the statement was run by a database analysis user.
	  • Space SQL - the statement was run by a database user of a space.
	  • Business Layer Modeling - the statement was run in the *Business Builder*.
	  • Data Layer Modeling - the statement was run in the data preview of the view editor in the *Data Builder*.
	  • DWC Space Management - the statement was run in the *Space Management*, for example, when deploying an object.
	  • DB Usergroup - the statement was run by a user of a database user group.
	  • DWC Administration - the statement was run for an administration task such as writing a task framework status.
	  • System - any other SAP HANA system statement.

Activity	Shows the action that was performed. For example: update, compile, select.
Object Name	If the statement is related to a task, it shows the name of the object for which the statement was run.
Schema Name	Shows the name of the schema in which the statement is run.
SAP HANA Peak Memory	Shows the maximum amount of memory (in MiB) the statement has used during the runtime in SAP HANA. **Note** You can see the information if the option *Enable Expensive Statement Tracing* is enabled and if the statement exceeds the thresholds specified in *(Configuration)* → *Monitoring*. Otherwise, no number is displayed.

Property	Description
SAP HANA CPU Time	Shows the amount of CPU time (in ms) the statement has used in SAP HANA. **Note** You can see the information if the option *Enable Expensive Statement Tracing* is enabled and if the statement exceeds the thresholds specified in *(Configuration)* → *Monitoring*. Otherwise, no number is displayed. **Note** The CPU time indicates how much time is used for all threads. It means that if the CPU time is significantly higher than the duration of the statement, then many threads are used. If many threads are used for a long time, no other tasks should be scheduled at that point in time, or resource bottlenecks may occur and tasks may even be canceled.
Statement Details	Shows the *More* link that you can click to view the complete SQL statement. **Note** For MDS queries - If you've enabled the tracing of MDS information (see Analyze Monitoring Data in a Space [page 197]), the payload of the MDS query that is run by SAP Analytics Cloud is displayed. If identified in the payload, the following information is also displayed: story ID, story name and data sources. You can copy or download the displayed information.
Parameters	Shows the values of the parameters of the statement that are indicated by the character "?" in the popup that opens when clicking *More* in the *Statement Details* column.
Out-of-memory	Shows if the statement has an out-of-memory error ("Yes" is then displayed) or not ("No" is then displayed).
Task Log ID	If the statement is related to a task, it shows the identifier of the task within a link, which takes you to the *Tasks* tab filtered on this task.
Error Code	If the statement has failed, it shows the numeric code of the SQL error. See SQL Error Codes in the *SAP HANA SQL Reference Guide for SAP HANA Platform*.
Error Message	If the statement has failed, it shows a description of the SQL error.
Workload Class	If the statement has an out-of-memory error, it shows the name of the workload class whose limit has been exceeded.

Note

Data on statements are kept for a time that depends on the thresholds specified in *(Configuration)* → *Monitoring*. As a certain number of statements are kept (30.000 by default), if very low thresholds are set, the time period may be very low (for example, only a few hours). To keep the statements for a longer time, the thresholds should be set accordingly.

Show/Hide, Filter, Sort and Reorder Task and Statement Columns

You can control the tables in *Tasks* and *Statements* in the following ways:

- Reorder the columns by drag and drop.
- Sort on a column by clicking the column header and then clicking *(Sort Ascending)* or *(Sort Descending)*.
- Filter on a column by clicking the column header, then clicking *(Filter)*
  The *Define Filter* dialog opens and advanced filtering options are available:
  1. Chose the appropriate section for your filter. If your filter is meant to include data in the table (you could say "I want my Data Preview to show"), add your filter in the *Include* section. If your filter is meant to exclude data from the table (you could say "I want my Data Preview to hide"), add your filter in the *Exclude* section. When in the appropriate section, click *(Add Filter)* to add a filter.
  2. Select a column to filter on, a filtering option, and a value. You can add several filters. Click *OK* to apply the filter(s). The currently applied filters are displayed above the table.

     > **Example**
     >
     > To only see the tasks that have failed on remote tables, in the *Include* area, select the column *Object Type*, then the filtering value *contains* and enter "REMOTE". Then, add a filter, select the column *Status*, then the filtering value *contains* and enter "FAILED". Once applied, the filter is displayed above the table.

     > **Note**
     >
     > - The filtering options available depend on the data type of the column you filter on.
     > - Filters applied to text columns are case-sensitive.
     > - You can enter filter or sort values in multiple columns.

     > **Note**
     >
     > If you filter on one of the following columns and you enter a number, use the "." (period) character as the decimal separator, regardless of the decimal separator used in the number formatting that you've chosen in the general user settings (*Settings* > *Language & Region*): *SAP HANA Peak Memory*, *SAP HANA CPU Time*, *SAP HANA Used Memory* and *SAP HANA Used Disk*.
  3. Click *Clear Filter* in the filter strip or *(Remove Filter)*in the *Define Filter* dialog to remove the filter.
- Show or hide columns by clicking *(Columns Settings)* to open the *Columns Settings* dialog, selecting columns as appropriate. To return to the default preview columns, click *Reset*.
- Refresh the table at any time by clicking *Refresh*.

8.1 Analyze Monitoring Data in a Space

Define the two spaces dedicated to monitoring SAP Datasphere (such as monitoring the database for resource consumption).

Prerequisites

You need the DW Administrator role to access the *Monitoring* page and select a space.

Context

Monitoring information provided by monitoring views can be sensitive as it includes information on all spaces and views. This is why these views are not accessible to all users by default in SAP Datasphere. However, as an administrator, you can select two spaces dedicated to monitoring information. The users who are assigned to the spaces and have modeling privileges, can then access the monitoring views in the Data Builder.

As the monitoring spaces you choose will provide unfiltered access to monitoring views, be aware that the users assigned to the spaces will be able to see all metadata and object definitions of all spaces.

You can dedicate one or two of these spaces to monitoring:

- Space to choose among all spaces - Choose a space that you want to dedicate to monitoring views.

  > ⓘ Note
  >
  > If you have already selected a space dedicated to monitoring before version 2021.19, you need to select another space, then select the initial space again so that you can access all the views.

- `<SAP_ADMIN>` space - This space is dedicated to the pre-configured monitoring views provided as business content by SAP via the *Content Network*. First create the space with the space ID `<SAP_ADMIN>` and the space name `<Administration (SAP)>`, enable access to it and import the package from the *Content Network*.

  > ⓘ Note
  >
  > Please do not create a space with the space ID `<SAP_ADMIN>` for another purpose.

  For more information about the package prepared by SAP, see Monitoring Tasks, Logs and Schedules With Dedicated Monitoring Views [page 199].

Monitoring Views Available in Monitoring Spaces

Once you've chosen one or two monitoring spaces, the users assigned to the spaces can access the following monitoring views in the *Data Builder* editors:

- **SAP HANA SYS Schema Monitoring Views**
  All SAP HANA monitoring views start with M_. For more information about all the monitoring views available, see Monitoring Views in the *SAP HANA Cloud, SAP HANA Database SQL Reference Guide*.
  The views for expensive statements are M_EXPENSIVE_STATEMENTS and M_EXPENSIVE_STATEMENT_EXECUTION_LOCATION_STATISTICS. For more information, see M_EXPENSIVE_STATEMENTS and M_EXPENSIVE_STATEMENT_EXECUTION_LOCATION_STATISTICS in the *SAP HANA Cloud, SAP HANA Database SQL Reference Guide*.
  The view M_MULTIDIMENSIONAL_STATEMENT_STATISTICS provides extensive information about MDS queries. For more information, see M_MULTIDIMENSIONAL_STATEMENT_STATISTICS System View in the *SAP HANA Cloud, SAP HANA Database SQL Reference Guide*.
- **SAP HANA _SYS_STATISTICS Schema Statistics Service Views**
  For more information, see Embedded Statistics Service Views (_SYS_STATISTICS schema).
- **SAP HANA _SYS_BI Schema Tables and Views**
  For more information, see BIMC Tables and Views in the *SAP HANA Cloud, SAP HANA Analytics Catalog (BIMC Views) Reference*.
- **SAP HANA DWC_GLOBAL Schema Monitoring Views**
  For more information, see Monitoring Tasks, Logs and Schedules With Dedicated Monitoring Views [page 199].
- **SAP Datasphere Monitoring Views (Delivered via the Content Network)**
  These views are available in the `<SAP_ADMIN>` space if you've enabled the space and imported the dedicated monitoring package from the *Content Network*.

Expensive Statement Tracing

In the area *Expensive Statement Tracing*, you can enable expensive statement tracing to analyze individual SQL queries whose execution exceeds one or more thresholds that you specify.

- The information on statements that exceed the specified thresholds are included in dedicated views that you can access in the selected monitoring space.
- All the database statements that exceed the thresholds specified for memory consumption (*Threshold Memory*) and runtime (*Threshold Duration*) are displayed in the widgets and tables of the *System Monitor*. If expensive statement tracing is not enabled, then statement information and errors are not traced and you cannot see them in the *System Monitor* (see Monitoring SAP Datasphere [page 188]).
- By default, 30 000 records maximum are stored in the monitoring tables. You can change this number, which will impact the views dedicated to monitoring and information related to statements in the *System Monitor*. For example, if about 5 days are traced in the expensive statement tables and you don't want to change the thresholds, you can double the number of records in *In-Memory Tracing Records* so that about 10 days are traced. Be aware that increasing this number will also increase the used storage.

MDS Information Tracing

To analyze individual SAP HANA multi-dimensional services (MDS) query, you can enable the tracing of MDS information in the area *MDS Information Tracing*.

If the tracing is enabled, you can view information on MDS queries when clicking *More* in the column *Statement Details* of the *Statements* tab in the *System Monitor* (see Monitoring SAP Datasphere [page 188]).

Procedure

1. Go to *(Configuration)* → *Monitoring*.
2. Select a space from the drop-down list and click *Confirm Selected Space*.
3. If you've created the `<SAP_ADMIN>` space and you want to enable it, click *Enable access to SAP Monitoring Content Space*. If there isn't any space named `<SAP_ADMIN>` in your tenant, this is not available for selection.
4. To trace expensive statements, select *Enable Expensive Statement Tracing*, specify the following parameters to configure and filter the trace details, then save your changes.
   - *In-Memory Tracing Records*: Change, if needed, the maximum number of records that are stored in the monitoring tables.
   - *Threshold CPU Time*: Specifies the threshold CPU time of statement execution. When set to 0, all SQL statements are traced.
   - *Threshold Memory*: Specifies the threshold memory usage of statement execution. When set to 0, all SQL statements are traced.
   - *Threshold Duration*: Specifies the threshold execution time. When set to 0, all SQL statements are traced.
   - *Trace Parameter Values*: In SQL statements, field values may be specified as parameters (using a "?" in the syntax). If these parameter values are not required, then do not select the option to reduce the amount of data traced.

   For more information about these parameters, see Expensive Statements Trace in the *SAP HANA Cloud, SAP HANA Database Administration Guide*.
5. To trace MDS information, select *Enable MDS Information Tracing* and save.

Troubleshooting

After the configuration, if you face authorization issues due to insufficient privilege on the monitoring views (suffix V_EXT) in your monitoring space, the solution is to choose another space as the monitoring space in the configuration UI, and then select the existing current monitoring space again.

8.1.1 Monitoring Tasks, Logs and Schedules With Dedicated Monitoring Views

Monitor tasks and schedules execution across spaces using monitoring views.

As an administrator, you need to keep an eye on how tasks and schedules are running across spaces. Gathering information from different logs might be time consuming and you need to find an easy way to collect this information.

> **Note**
>
> You can also monitor all tasks that are run across all spaces directly in the *System Monitor*. For more information, see Monitoring SAP Datasphere [page 188].

Using SAP HANA Monitoring Views From the DWC_GLOBAL Schema

The following monitoring views are ready to use in the DWC_GLOBAL schema and can be recognized as they have a naming convention Suffix_V_EXT:

- SPACE_SCHEMAS_V_EXT:

Column	Function
SPACE_ID	Identifier of the SAP Datasphere space. Note that one space can contain several schemas.
SCHEMA_NAME	Name of the schema used to run the task.

- SPACE_USERS_V_EXT:

Column	Function
SPACE_ID	Identifier of the SAP Datasphere space. Note that one space can contain several users.
USER_NAME	Identifier of the user.
USER_TYPE	Type of user, such as space technical user (for example database user for open SQL schemas) or global user.

- TASK_SCHEDULES_V_EXT:

Column	Key	Function	Values
SPACE_ID	X	Identifier of the SAP Datasphere space which contains the object with the defined schedule.	
OBJECT_ID	X	Identifier of the SAP Datasphere object for which the schedule is defined.	
APPLICATION_ID	X	Identifier of the type of object	For example: VIEWS, REMOTE_TABLES, DATA_FLOWS, TASK_CHAINS
ACTIVITY	X	Identifier of the type of activity applied to the object. **Note** For each application, you can have multiple activities (for example, replicating or deleting data)	For example: PERSIST (View), EXECUTE (Dataflow), REPLICATE (Remote Tables), RUN_CHAIN (Task Chain)
OWNER		Identifier of the responsible of the schedule, schedule executed on users behalf, consent is checked against	< DWC User ID >
CRON		Defines the recurrence of a schedule in CRON format	NULL (no schedule defined, or a SIMPLE schedule is defined) For example: "0 */1 * * *" for hourly. For

Column	Key	Function	Values
			more information on CRON schedule, see Schedule a Data Integration Task (with Cron Expression)
FREQUENCY		Defines the recurrence of a schedule in json format (simple format)	NULL (no schedule defined, or a CRON schedule is defined) or schedule definition, for example Daily + start date + time + duration. For more information, see Schedule a Data Integration Task (Simple Schedule)
CHANGED_BY		User who last changed the schedule configuration.	
CHANGED_AT		Timestamp containing Date and Time, at which the schedule was last changed.	

- TASK_LOGS_V_EXT

Column	Key	Function	Values
TASK_LOG_ID	X	Uniquely identifies an execution of a task.	
SPACE_ID		Identifier of the SAP Datasphere space which contains the object with the defined schedule.	
APPLICATION_ID		Identifier of the type of object	For example: VIEWS, REMOTE_TABLES, DATA_FLOWS, TASK_CHAINS
OBJECT_ID		Identifier of the SAP Datasphere object for which the schedule is defined.	
ACTIVITY		For each application there could be multiple activities, e.g. replicating or deleting data	For example: PERSIST (View), EXECUTE (Dataflow), REPLICATE (Remote Tables), RUN_CHAIN (Task Chain)
PEAK_MEMORY		Captures the highest peak memory consumption (in bytes). Note: this value is not available for all apps. To capture this value when it's available, you must check *Enable Expensive Statement Tracing* in *Configuration* > *Monitoring*. For more information, see Monitoring SAP Datasphere [page 188].	Null (not available for the application, or not measured because the *Enable Expensive Statement Tracing* is not set, or the threshold defined is not reached), 0 or value of the memory consumption.

Column	Key	Function	Values
PEAK_CPU		Total CPU time (in microseconds) consumed by the task. Note: This value is not available for all apps.To capture this value when it's available, you must check *Enable Expensive Statement Tracing* in *Configuration* > *Monitoring*. For more information, see Monitoring SAP Datasphere [page 188].	Null (not available for the application, or not measured because the *Enable Expensive Statement Tracing* is not set, or the threshold defined is not reached), 0 or value of the CPU time consumption.
RECORDS		Shows the number of records of the target table after the task has finished running.	Null (not applicable or not measured), 0 or number of records.
START_TIME		Timestamp containing Date and Time, at which the scheduled task was started.	
END_TIME		Timestamp containing Date and Time, at which the scheduled task was stopped.	
STATUS		Reports if this task execution is still running, completed or failed.	
TRIGGERED_TYPE		Indicates if task execution was triggered manually (DIRECT) or via schedule (SCHEDULED)	
APPLICATION_USER		The user on whose behalf the schedule was executed (the owner at this point in time)	
DURATION		Duration of the task execution (also works for ongoing execution)	
START_DATE		Date when the scheduled task was started.	

- TASK_LOG_MESSAGES_V_EXT

Column	Key	Function
TASK_LOG_ID	X	Uniquely identifies an execution of a task
MESSAGE_NO	X	Order sequence of all messages belonging to a certain Tasklog ID.
SEVERITY		Indicates if the message provides general information (INFO) or error information (ERROR)
TEXT		The message itself
DETAILS		Technical additional information. For example, it can be an error stack or a correlation ID.

- TASK_LOCKS_V_EXT

Column	Key	Function
LOCK_KEY	X	Identifier, flexible field as part of the lock identifier, usually set to WRITE or EXECUTE.
APPLICATION_ID	X	Identifier of the type of object.

Column	Key	Function
SPACE_ID	X	Identifier of the SAP Datasphere space which contains the object with the defined schedule.
OBJECT_ID	X	Identifier of the SAP Datasphere object for which the schedule is defined.
TASK_LOG_ID		Uniquely identifies the task execution that set the lock.
CREATION_TIME		Indicates when the lock has been set.

> **Note**
>
> Cross-space sharing is active for all SAP HANA monitoring views. The row level access of shared views is bound to the space read access privileges of the user who consumes the view.

You can then choose a space that is dedicated to task framework monitoring. For more information, see Analyze Monitoring Data in a Space [page 197].

Using SAP Datasphere Monitoring Views (Delivered via the Content Network)

Prerequisites:

- You need to create a space with the space ID `<SAP_ADMIN>` and the space name `<Administration (SAP)>` and to configure it as a monitoring space by enabling the toggle *Enable Access to SAP Monitoring Content Space*:

  

  For more information, see Analyze Monitoring Data in a Space [page 197].
- You need to import the package that contains monitoring views via the *Content Network*. Import the package `<Technical Content: Task Monitoring>`. For more information, see Importing SAP and Partner Business Content from the Content Network.

Thanks to SAP Datasphere pre-configured monitoring views, you can monitor data integration tasks in a more flexible way. They are based on top of the V_EXT HANA views (see previous section), and we enriched them with further information as preparation for consumption in an SAP Analytics Cloud Story.

For more information about how to create such a story,

See the blogs SAP Datasphere: Data Integration Monitoring – Sample Content for Reporting (published in October 2021) and SAP Datasphere: Data Integration Monitoring – Running Task Overview (published in November 2021).

For more information about data integration, see Importing SAP and Partner Business Content from the Content Network.

> **Restriction**
>
> Content modification is not recommended as any update will overwrite custom changes.

The following views are available:

- SAP_TCT_TASK_LOGS_V_R_01: Monitoring: Task Execution Headers
  - **The view exposes**:
    - Task properties, such as duration and execution status (e.g. failed, completed, ...).
    - Various measures for counting tasks (e.g. failed).
    - The schedule description.
    - Locking status
  - It uses the views TASK_LOCKS_V_EXT, TASK_SCHEDULES_V_EXT and TASK_LOGS_V_EXT:

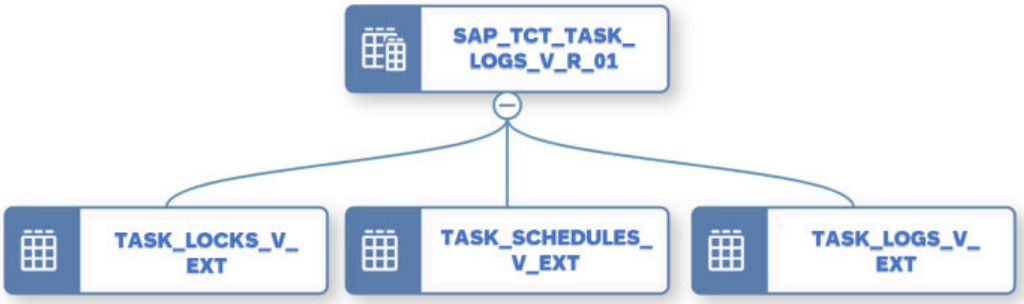

  - Best Practice: To enable the navigation between SAP Datasphere and SAP Analytics Cloud, you must change the constant for the url_host to your SAP Datasphere instance. Open the view in the view editor, and update the URL host:

```
/* -- Enable jump from SAC report to DWC detailed view
-- URL Pattern: 'https://(URL_HOST)(URL_PATH)/(OBJECT_ID)' */
'<my-dwc-system>.hcs.cloud.sap' URL_HOST,
concat(concat('/dwaas-ui/index.html#/dataintegration&/di/', base."SPACE_ID"), base."MONITOR_TYPE") URL_PATH
```

- SAP_TCT_TASK_SCHEDULE_V_R_01: Monitoring: Schedule Properties
  - **The view exposes** the properties of a data integration schedule.

- It uses the view TASK_SCHEDULES_V_EXT and adds a row-count to be compatible with OLAP reporting:

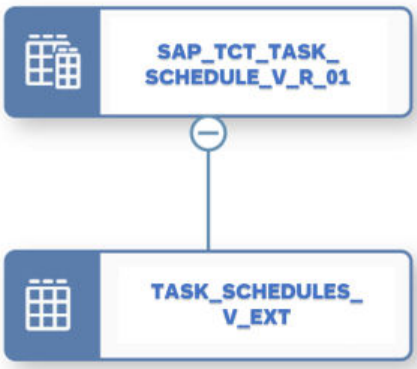

- SAP_TCT_TASK_MSG_V_R_01: Monitoring: Task Execution Items
  - **The view exposes**:
    - All messages occurring during data integration monitoring.
    - Error code, header line and first stack line parsed out from detailed message.
    - An indicator that the task_id has an error (facilitate filtering of messages).
  - It uses the views TASK_LOG_MESSAGES_V_EXT and TASK_LOGS_V_EXT:

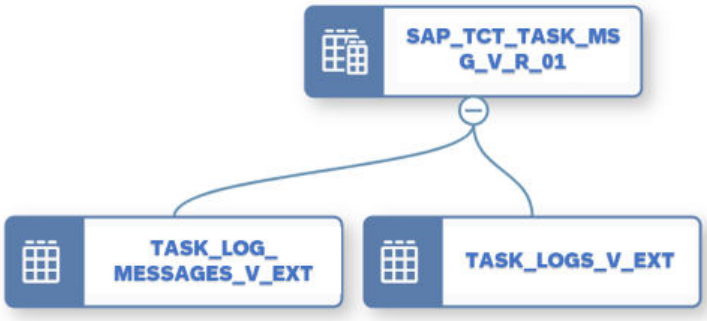

- - Best Practice: To enable the navigation between SAP Datasphere and SAP Analytics Cloud, you must change the constant for the url_host to your SAP

Datasphere instance. Open the view in the view editor, and update the URL host:

```
/* -- Enable jump from SAC report to DWC detailed view
-- URL Pattern: 'https://(URL_HOST)(URL_PATH)/(OBJECT_ID)' */
'<my-dwc-system>.hcs.cloud.sap' URL_HOST,
concat(concat('/dwaas-ui/index.html#/dataintegration&/di/', base."SPACE_ID"), base."MONITOR_TYPE") URL_PATH
```

8.2 Monitor Database Operations with Audit Logs

Monitor the read and change actions (policies) performed in the database with audit logs, and see who did what and when.

If Space Administrators have enabled audit logs to be created for their space (see Enable Audit Logging), you can get an overview of these audit logs. You can do analytics on audit logs by assigning the audit views to a dedicated space and then work with them in a view in the Data Builder.

> **Note**
>
> Audit logs can consume a large quantity of GB of disk in your database, especially when combined with long retention periods (which are defined at the space level). You can delete audit logs when needed, which will free up disk space. For more information, see Delete Audit Logs [page 206].

1. Choose a space that will contain the audit logs.

   Go to *System* > *Configuration* > *Audit*. Enable to save and later display the audit logs directly in a certain space by choosing a space from the drop-down list. We recommend to create a dedicated space for audit logs, as you might not want all users to view sensitive data.
2. Open the *Data Builder*, create a view, and add one or more of the following views from the DWC_AUDIT_READER schema as sources:
   - `DPP_AUDIT_LOG` - Contains audit log entries.
   - `AUDIT_LOG_OVERVIEW` - Contains audit policies (read or change operations) and the number of audit log entries.
   - `ANALYSIS_AUDIT_LOG` - Contains audit log entries for database analysis users. For more information, see Create a Database Analysis User to Debug Database Issues [page 208].

8.2.1 Delete Audit Logs

Delete audit logs and free up disk space.

All spaces for which auditing is enabled, are listed in the *Audit Log Deletion* area.

For each space, you can delete separately all the audit log entries recorded for read operations and all the audit log entries recorded for change operations. All the entries recorded before the date and time you specify are deleted.

1. Go to *System* > *Configuration* > *Audit*.
2. Select the spaces and the audit policy names (read or change) for which you want to delete all audit log entries and click *Delete*.

3. Select a date and time and click *Delete*.
   All entries that have been recorded before this date and time are deleted.
   Deleting audit logs frees up disk space, which you can see in the *Used Disk* bar of the *Space Management* page.

> **Note**
>
> Audit logs are automatically deleted when performing the following actions: deleting a space, deleting a database user (open SQL schema), disabling an audit policy for a space, disabling an audit policy for a database user (open SQL schema), unassigning an HDI container from a space. Before performing any of these actions, you may want to export the audit log entries, for example by using SAP HANA Database Explorer (see Export Audit Logs).

8.3 Monitor Object Changes with Activities

Monitor the changes that users perform on modeling objects (such as spaces and tables) as well as changes to the system configuration (such as roles and users).

Actions that are performed by users are logged in *Security* > *Activities*.

For example:

- Space creation and changes
- Table changes
- Role changes and assignments
- Logged users

On the *Activities* page, you can:

- View the activities and filter on specific activities.
- Download the activity log for a specific time period.
- Delete the activity log for a specific time period.

> **Note**
>
> To delete the activity log, you must be granted the privilege *Activity Log* with the permission *Delete*, which is included in the system owner role and which you can include in a custom role.

For more information, see Track User Activities in the *SAP Analytics Cloud Help*.

8.4 Create a Database Analysis User to Debug Database Issues

A database analysis user is an SAP HANA Cloud database user with wide-ranging privileges. It can be used to support monitoring, analyzing, tracing, and debugging of your SAP Datasphere run-time database.

Context

A user with the *DW Administrator* role can create a database analysis user.

> **Note**
>
> You should only create a database analysis user to resolve a specific database issue and then delete it immediately after the issue is resolved. This user can access all SAP HANA Cloud monitoring views and all SAP Datasphere data in all spaces, including any sensitive data stored there.

Procedure

1. In the side navigation area, click (*System*) > (*Configuration*) > *Database Access* > *Database Analysis Users*.
2. Click *Create* and enter the following properties in the dialog:

Property	Description
Database Analysis User Name Suffix	Enter the suffix, which is used to create the full name of the user. Can contain a maximum of 31 uppercase letters or numbers and must not contain spaces or special characters other than _ (underscore). See Rules for Technical Names [page 106].
Enable Space Schema Access	Select only if you need to grant the user access to space data.
Database analysis user expires in	Select the number of days after which the user will be deactivated. We strongly recommend creating this user with an automatic expiration date.

3. Click *Create* to create the user.

   The host name and port, as well as the user password are displayed. Note these for later use.
4. Select your user in the list and then click one of the following and enter your credentials:

   - *Open SAP HANA Cockpit* - Open the *Database Overview* > *Monitoring* page for the SAP Datasphere run-time database, which offers various monitoring tools.
     For more information, see Using the Database Overview Page to Manage a Database).
   - *Open Database Explorer* - Open an SQL Console for the SAP Datasphere run-time database.
     For more information, see Getting Started With the SAP HANA Database Explorer).
     A database analysis user can run a procedure in Database Explorer to stop running statements. For more information, see Stop a Running Statement [page 209].

> **Note**
>
> All actions of the database analysis user are logged in the `ANALYSIS_AUDIT_LOG` view, which is stored in the space that has been assigned to store audit logs (see Enable Audit Logging).
>
> The audit logs entries are kept for 180 days, after which they are deleted.

8.4.1 Stop a Running Statement

Using a database analysis user, you can stop a statement that is currently running.

You may for example want to stop a statement that has been running for a long time and is causing performance issues.

You can only stop a statement that has been run by space users, analysis users, user group users and Data Provisioning Agent users.

In SAP HANA Database Explorer, run a database procedure using the following syntax:

```
CALL "DWC_GLOBAL"."STOP_RUNNING_STATEMENT"('<ACTION>', '<CONNECTION_ID>')
```

Complete the parameters as follows:

Parameter	Value	Description
ACTION	CANCEL	Enter CANCEL to run the statement ALTER SYSTEM CANCEL [WORK IN] SESSION (see ALTER SYSTEM CANCEL [WORK IN] SESSION Statement (System Management) in the *SAP HANA Cloud, SAP HANA Database SQL Reference Guide*.)
	DISCONNECT	Enter DISCONNECT to run the statement ALTER SYSTEM DISCONNECT SESSION (see ALTER SYSTEM DISCONNECT SESSION Statement (System Management) in the *SAP HANA Cloud, SAP HANA Database SQL Reference Guide*.)
CONNECTION_ID		Enter the ID of the connection to the database, which corresponds to the statement that you want to stop.

For more information on database explorer, Getting Started With the SAP HANA Database Explorer.

8.4.2 Delete a Database Analysis User

Delete your database analysis user as soon as the support task is completed to avoid misuse of sensitive data.

Procedure

1. In the side navigation area, click (*System*) > (*Configuration*) > *Database Access* > *Database Analysis Users*.
2. Select the user you want to delete and then click *Delete*.

8.5 Configure Notifications

Configure notifications about system events and network connection issues, and define the SMTP server to be used for email deliveries.

Notify All Users about Network Connection Issues

When there are problems with a system, your users would like to know whether it is something that they control or if the issues are related to the network. You can't create messages for all situations, but you can let them know when the network connection is unstable.

To turn on the connection notification:

1. In the side navigation area, click *System* > *Administration* > *Notifications*.
2. To enable editing of all settings on the page, click *Edit*.
3. In the *Connections Notifications* section, change the toggle to *ON*.
4. Click *Save* to commit your changes.

When the notification is on, everyone who uses the application on that tenant will see the notification in the top right corner of their application.

Configure Custom SMTP Server

Configuring an email server of your choice ensures greater security and flexibility while delivering email for your business.

1. In the side navigation area, click *System* > *Administration* > *Notifications*.
2. To enable editing of all settings on the page, click *Edit*.
3. In the *Email Server Configuration* section, select *Custom*, and complete the following properties.

4. Click *Check Configuration* or *Save* to successfully validate the configuration details.

8.6 Check Consent Expirations

View a list of users whose authorization consent will expire in less than four weeks.

To view a list of users whose authorization consent will expire within the next four weeks, click (*Configuration*) ▶ *Tasks*. Then, in the *Consent Expiration* section of the *Tasks* page, click the *View Expiration List* link. SAP Datasphere now displays a dialog in which you can view a list of users whose authorization consent will expire within a given timeframe.

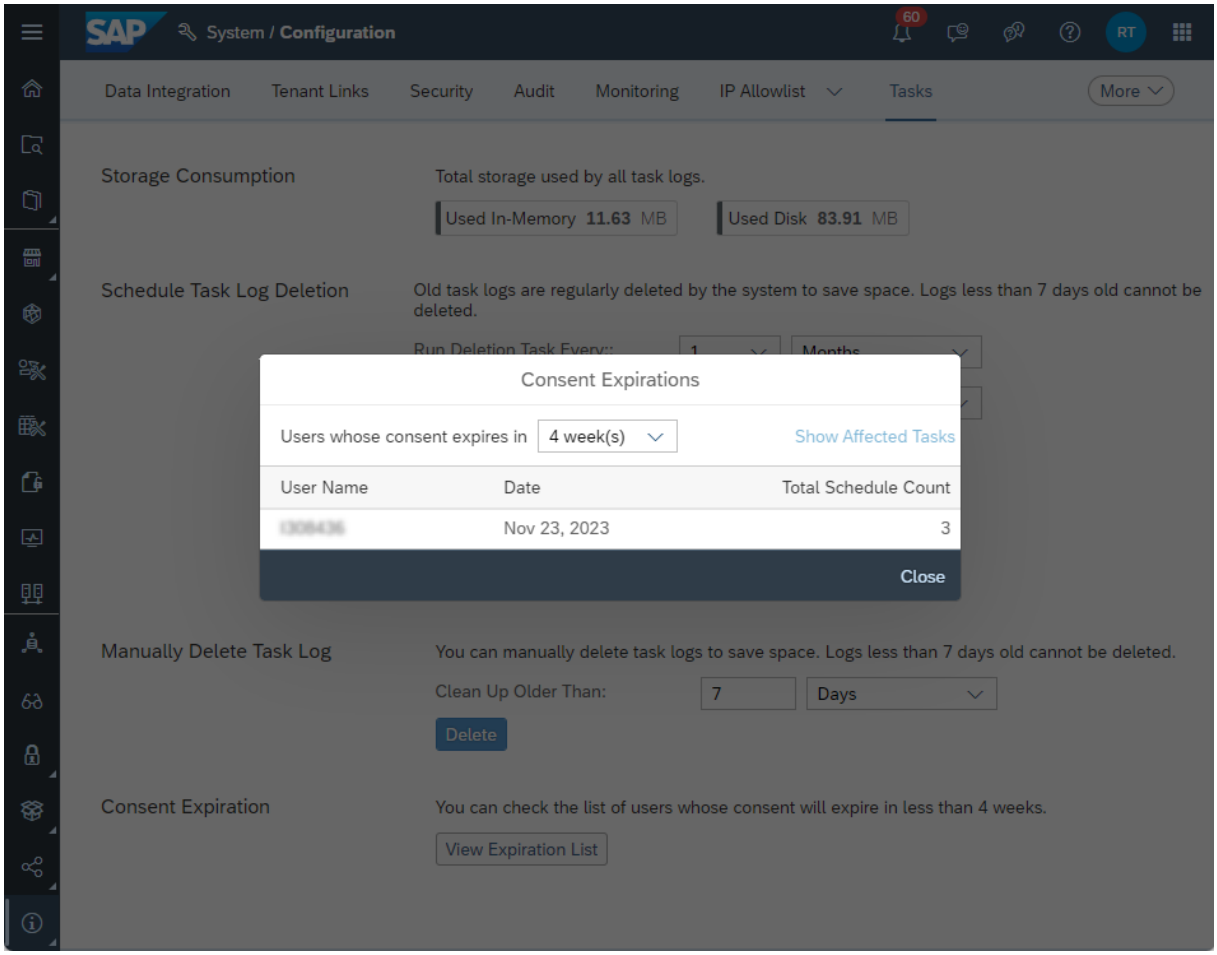

By default, the dialog displays a list of users whose consent will expire within four weeks. You can change the default expiration timeframe to anywhere between one and four weeks. In addition to displaying the list of users whose consent will soon expire, you can also select a user in the list and click the *Show Affected Tasks* link to view the collection of tasks that user has scheduled.

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 an 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.

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.

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.

Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.

www.sap.com/contactsap

THE BEST RUN SAP